SIAM J. Financial Math.
Vol. 11, No. 2, pp. 470–493

Optimal Execution with Rough Path Signatures*

Jasdeep Kalsi†, Terry Lyons‡, and Imanol Perez Arribas§

Abstract. 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 optimization 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 closed-form solutions of the optimal trading speed are not known. In the latter case, we obtain favorable comparisons with standard execution strategies.

Key words. optimal execution, rough path theory, signatures, high-frequency trading

AMS subject classifications. 91G80, 91G60, 60G99

DOI. 10.1137/19M1259778

1. Introduction.

1.1. Overview. The problem of optimal execution has attracted much interest following the original work on the problem by Bertsimas and Lo in [5] and Almgren and Chriss in [1]. 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 optimization. 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 could result in "walking the book" and accepting unfavorable 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

*Received by the editors May 3, 2019; accepted for publication (in revised form) February 26, 2020; published electronically April 27, 2020. Opinions expressed in this paper are those of the authors and do not necessarily reflect the view of JP Morgan.
https://doi.org/10.1137/19M1259778
Funding: This work was supported by The Alan Turing Institute under EPSRC grant EP/N510129/1.
†Mathematical Institute, University of Oxford, Oxford, OX2 6GG, UK (Jasdeep.Kalsi@maths.ox.ac.uk).
‡Mathematical Institute, University of Oxford, and The Alan Turing Institute (London), UK (Terry.Lyons@maths.ox.ac.uk).
§Mathematical Institute, University of Oxford, J.P. Morgan (London), and The Alan Turing Institute (London), UK (Imanol.Perez@maths.ox.ac.uk).

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 objective 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.

1.2. Paper outline. The aim of this paper 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 depending 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 [12, 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 [26] and Cartea and Jaimungal in [11]. In [26], 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 [11], 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 affected by trades in each other. Such a setting is presented in the article by Mastromatteo et al. [29].

We begin the paper by introducing in section 2 the framework we will be working on. 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 3 is dedicated to calculating approximate solutions to the control problem. We first reformulate the problem in terms of the signature, and then we approximate the optimal trading speed by a finite-dimensional, computationally tractable optimization problem. In section 4, we provide examples of interesting extensions of the approach as it was presented in sections 2 and 3, such as the multiple asset problem which appears in [29], and more exotic models where additional multidimensional noise is assumed to provide exogenous information about the price dynamics. Finally in sections 5 and 6, we provide numerical evidence of the performance of the model. Good approximations to the optimal strategies in the settings [12, section 6.5], [26] and [11] are obtained, and we also investigate the problem in the case where the underlying price process is a fractional Brownian motion. Moreover, we explore in section 6 the effectiveness of our methodology on real market data.

2. Framework.

2.1. The market. Let $X : [0,T] \to \mathbb{R}$ with $X_0 = 1$ be a continuous stochastic process, which denotes 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. We also denote by $\widehat{X} : [0,T] \to \mathbb{R}^2$ the 2-dimensional process $\widehat{X}_t := (t, X_t)$. We impose the following technical assumption on the midprice process.

Assumption 2.1. *The unaffected midprice process $\widehat{X}$ admits a canonical geometric rough path lift almost surely (see Appendix A).*

The authors refer the reader to Appendix A for an introduction to rough paths or [28, 23] for a more thorough overview of rough path theory. This technical assumption is not restrictive—on the contrary, our market framework is very general in the sense that it includes most of the examples considered in the literature. In particular, our framework includes the following.

1. *Continuous semimartingales.* In the literature [9, 11, 10, 26], the midprice process is often modelled as a continuous semimartingale. Continuous semimartingales can be lifted to geometric rough paths [27, 16, 23], and therefore they fit into our framework.
2. *Fractional Brownian motion.* Our framework also includes the setting where the midprice is modelled by a fractional Brownian motion with Hurst parameter $H \geq 1/4$, see [17].
3. *Tick data.* In [21], 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.

2.2. Signature of the price process. Throughout this paper, we use the following notation for *multi-indices.*

Definition 2.2. *We denote the set of all multi-indices of length at most $N \in \mathbb{N}$ by*

$$\mathcal{I}_N := \{(i_1, i_2, \dots, i_n) \,:\, i_1, \dots, i_n \in \{1,2\}, n = 0, \dots, N\}.$$

Similarly, we define the set of all multi-indices of arbitrary length by $\mathcal{I} := \bigcup_{N=1}^{\infty} \mathcal{I}_N$. Given a multi-index $\mathbf{w} = (i_1, \dots, i_n) \in \mathcal{I}$, we denote by $|\mathbf{w}| := n$ its length. Given two multi-indices $\mathbf{w} = (i_1, \dots, i_n), \mathbf{v} = (j_1, \dots, j_m) \in \mathcal{I}$ we denote by $\mathbf{w} \sqcup \mathbf{v} := (i_1, \dots, i_n, j_1, \dots, j_m) \in \mathcal{I}$ their concatenation.

We now define a crucial object for this paper: the signature of the price process.

Definition 2.3 (signature of the unaffected midprice process). *Let $X : [0,T] \to \mathbb{R}$ be the unaffected midprice process, and write the 2-dimensional path $\widehat{X}_t = (\widehat{X}^1_t, \widehat{X}^2_t) := (t, X_t)$. The signature of $\widehat{X}$ over the interval $[s,t] \subset [0,T]$ is defined as the collection*

$$\widehat{\mathbb{X}}^{<\infty}_{s,t} := \left(\widehat{\mathbb{X}}^{\mathbf{w}}_{s,t}\right)_{\mathbf{w} \in \mathcal{I}},$$

where $\widehat{\mathbb{X}}^{\mathbf{w}}_{s,t} \in \mathbb{R}$ is defined recursively as

$$\widehat{\mathbb{X}}^{\mathbf{w}}_{s,t} := 1 \qquad \textit{for } |\mathbf{w}| = 0,$$
$$\widehat{\mathbb{X}}^{\mathbf{w}}_{s,t} := \widehat{X}^i_t - \widehat{X}^i_s \qquad \textit{for } \mathbf{w} = (i), i = 1, 2,$$
$$\widehat{\mathbb{X}}^{\mathbf{w}}_{s,t} := \int_s^t \widehat{\mathbb{X}}^{\mathbf{v}}_{s,u} d\widehat{X}^i_u \qquad \textit{for } \mathbf{w} = \mathbf{v} \sqcup (i), i = 1, 2. \tag{2.1}$$

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

$$\widehat{\mathbb{X}}^{\leq N}_{s,t} := \left(\widehat{\mathbb{X}}^{\mathbf{w}}_{s,t}\right)_{\mathbf{w}\in\mathcal{I}_N}.$$

We refer the reader to Appendix A for a formal definition of the signature and the integrals in (2.1). We note that there are numerically efficient approaches to compute signatures (see, for instance, [31]).

Definition 2.4 (linear functionals on signatures). *For each multi-index $\mathbf{w} \in \mathcal{I}$ define the linear functional $e_{\mathbf{w}}$ that acts on the signature by $\langle e_{\mathbf{w}}, \widehat{\mathbb{X}}^{<\infty}_{s,t}\rangle := \widehat{\mathbb{X}}^{\mathbf{w}}_{s,t}$. We also denote by 1^* the constant functional $\langle 1^*, \widehat{\mathbb{X}}^{<\infty}_{s,t}\rangle = 1$. These basic linear functionals are extended to linear functionals $\ell = \sum_{\mathbf{w}\in\mathcal{I}} \ell_{\mathbf{w}} e_{\mathbf{w}}$ with $\ell_{\mathbf{w}} \in \mathbb{R}$ such that all $\ell_{\mathbf{w}}$ but finitely many of them are zero. These linear functionals act on the signature by*

$$\langle \ell, \widehat{\mathbb{X}}^{<\infty}_{s,t}\rangle := \sum_{\mathbf{w}\in\mathcal{I}} \ell_{\mathbf{w}} \langle e_{\mathbf{w}}, \widehat{\mathbb{X}}^{<\infty}_{s,t}\rangle = \sum_{\mathbf{w}\in\mathcal{I}} \ell_{\mathbf{w}} \widehat{\mathbb{X}}^{\mathbf{w}}_{s,t}.$$

Notice that the sum is finite as only finitely many $\ell_{\mathbf{w}}$ are nonzero. Given a multi-index $\mathbf{v} \in \mathcal{I}$, we denote by $\ell \sqcup \mathbf{v}$ the linear functional

$$\ell \sqcup \mathbf{v} := \sum_{\mathbf{w}\in\mathcal{I}} \ell_{\mathbf{w}} e_{\mathbf{w}\sqcup\mathbf{v}}. \tag{2.2}$$

We define below the *shuffle product* of two multi-indices. Intuitively, the shuffle product accounts for all the possible ways of riffle shuffling two decks of cards—for multi-indices, it captures all possible ways of merging two multi-indices while preserving the relative order of each multi-index.

Definition 2.5 (shuffle product). *Given two multi-indices $\mathbf{w} = (i_1, \dots, i_n), \mathbf{v} = (j_1, \dots, j_m) \in \mathcal{I}$, the shuffle product of $e_{\mathbf{w}}$ and $e_{\mathbf{v}}$ is the linear functional denoted by $e_{\mathbf{v}} ⧢ e_{\mathbf{w}}$ which is recursively defined as*

$$e_{\mathbf{w}} ⧢ e_{\mathbf{v}} := (e_{(i_1,\dots,i_{n-1})} ⧢ e_{\mathbf{v}}) \sqcup (i_n) + (e_{\mathbf{w}} ⧢ e_{(j_1,\dots,j_{m-1})}) \sqcup (j_m) \qquad \textit{if } |\mathbf{w}| \neq 0, |\mathbf{v}| \neq 0,$$
$$e_{\mathbf{w}} ⧢ e_{\mathbf{v}} := e_{\mathbf{w}} \qquad \textit{if } |\mathbf{v}| = 0,$$
$$e_{\mathbf{w}} ⧢ e_{\mathbf{v}} := e_{\mathbf{v}} \qquad \textit{if } |\mathbf{w}| = 0.$$

Recall that $\sqcup$ denotes concatenation; see (2.2). The shuffle product is extended by bilinearity to linear functionals: given two linear functionals $f = \sum_{\mathbf{w}\in\mathcal{I}} f_{\mathbf{w}} e_{\mathbf{w}}$ and $g = \sum_{\mathbf{w}\in\mathcal{I}} g_{\mathbf{w}} e_{\mathbf{w}}$, their shuffle product is

$$f ⧢ g := \sum_{\mathbf{w},\mathbf{v}\in\mathcal{I}} f_{\mathbf{w}} g_{\mathbf{v}}\, e_{\mathbf{w}} ⧢ e_{\mathbf{v}}.$$

We also denote $f^{⧢ k} := \underbrace{f ⧢ f ⧢ \cdots ⧢ f}_{k}$ for $k \in \mathbb{N}$.

The shuffle product will play a key role in this paper due to the following lemma, known as the *shuffle product property*. It guarantees that the product of two linear functions on the signature is another linear function on the signature, which is given in terms of the shuffle product.

Lemma 2.6 (shuffle product property [28, Theorem 2.15]). *Let f, g be two linear functionals (Definition 2.4). Then, we have*

$$\langle f, \widehat{\mathbb{X}}^{<\infty}_{s,t}\rangle\langle g, \widehat{\mathbb{X}}^{<\infty}_{s,t}\rangle = \langle f \sqcup\mathchoice{\mkern-7mu}{\mkern-7mu}{\mkern-3.2mu}{\mkern-3.8mu}\sqcup g, \widehat{\mathbb{X}}^{<\infty}_{s,t}\rangle \quad \forall\ 0 \le s \le t \le T.$$

2.3. trading speeds.

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

Definition 2.7 (trading speeds). *The space of trading speeds $\mathcal{T}$ is defined as the space of all continuous adapted processes with respect to the filtration generated by X. We also define the space of signature trading speeds, denoted by $\mathcal{T}_{sig}$:*

$$\mathcal{T}_{sig} := \left\{\ell_{\mathbf{w_1}}\widehat{\mathbb{X}}^{\mathbf{w_1}}_{0,t} + \cdots + \ell_{\mathbf{w_n}}\widehat{\mathbb{X}}^{\mathbf{w_n}}_{0,t} \,:\, \mathbf{w_1}, \ldots, \mathbf{w_n} \in \mathcal{I}, n \in \mathbb{N}, t \in [0, T]\right\}.$$

Notice that $\mathcal{T}_{sig} \subset \mathcal{T}$.

Intuitively, the trader 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. This intuition is incorporated into the definition of trading speeds $\mathcal{T}$: the speed at which the trader will trade is an adapted process.

Signature trading speeds, on the other hand, are a special class of trading speeds. These signature trading speeds are adapted processes whose value at each time t is given by a (finite) linear combination of the signature of the unaffected midprice process up to time t.

The reason why we consider signature trading speeds $\mathcal{T}_{sig}$ instead of the larger class of trading speeds $\mathcal{T}$ is twofold. First, as we will see, optimizing over $\mathcal{T}_{sig}$ is computationally easier than optimizing over $\mathcal{T}$. Second, the space of signature trading speeds $\mathcal{T}_{sig} \subset \mathcal{T}$ is large—see Lemma B.3. This lemma states that there exists a *large* compact set (large in the sense that with very high probability the sample paths of the unaffected midprice process are on that compact set) such that on that compact set any trading speed can be approximated by signature trading speeds to arbitrary accuracy. Hence, if one wants to optimize a certain objective function over $\mathcal{T}$, it makes sense to optimize it over $\mathcal{T}_{sig}$ instead; this is the approach we follow in this paper.

2.4. Market impact.

When a trader buys or sells a traded asset, her trading activity 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 [3] and the references therein). In this section introduce the market impact model that will be used in this paper.

If the trader decides to follow a signature trading speed $\theta \in \mathcal{T}_{sig}$, the *execution price*—i.e., the price the trader receives—is given by

$$P^\theta_t := X_t - \langle g^\theta, \widehat{\mathbb{X}}^{<\infty}_{0,t}\rangle, \tag{2.3}$$

where g^θ is a linear functional that models the market impact and depends on θ.

Example 2.8. The definition of the market impact, far from being restrictive, includes many examples studied in the literature. Indeed, let $\ell = \sum_{\mathbf{w}\in\mathcal{I}} \ell_{\mathbf{w}} e_{\mathbf{w}}$ be a linear functional, and let $\theta_t := \langle \ell, \widehat{\mathbb{X}}^{<\infty}_{0,t}\rangle$ be a signature trading speed. 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, $\langle g^\ell, \widehat{\mathbb{X}}^{<\infty}_{0,t}\rangle = \lambda\theta_t$ is the linear temporary market impact studied in [9, 11, 26].
2. *Permanent market impact.* In [9, 11, 10], a permanent market impact given by $\int_0^t \theta_s ds$ is considered. Setting $g^\ell := \ell \sqcup (1)$, where the concatenation operation $\sqcup$ was defined in (2.2), we have

$$\langle g^\ell, \widehat{\mathbb{X}}^{<\infty}_{0,t}\rangle = \left\langle \sum_{\mathbf{w}\in\mathcal{I}} \ell_{\mathbf{w}} e_{\mathbf{w}\sqcup(1)}, \widehat{\mathbb{X}}^{<\infty}_{0,t}\right\rangle = \sum_{\mathbf{w}\in\mathcal{I}} \ell_{\mathbf{w}} \langle e_{\mathbf{w}\sqcup(1)}, \widehat{\mathbb{X}}^{<\infty}_{0,t}\rangle = \sum_{\mathbf{w}\in\mathcal{I}} \ell_{\mathbf{w}} \widehat{\mathbb{X}}^{\mathbf{w}\sqcup(1)}_{0,t}$$
$$= \sum_{\mathbf{w}\in\mathcal{I}} \ell_{\mathbf{w}} \int_0^t \widehat{\mathbb{X}}^{\mathbf{w}}_{0,s} ds = \int_0^t \langle \ell, \widehat{\mathbb{X}}^{<\infty}_{0,s}\rangle ds = \int_0^t \theta_s ds,$$

   so that the permanent market impact is included in our framework.
3. *Transient market impact.* In [25, 18, 19] the authors considered a transient market impact that is given by $\int_0^t K(t-s)\theta_s ds$, where $K(x) := \exp(-\rho x)$ for $\rho > 0$ constant. Then, we can find a linear functional g^ℓ such that

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

   to arbitrary accuracy. Appendix C explains how such an approximation can be derived.
4. More generally, market impacts modelled by functions of the form $G(\theta, X)$ are well-approximated by linear functions on the signature, and they are therefore included in our framework.

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 [3] 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 is 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. Next, we introduce the optimal execution problem that will be studied in this paper.

Definition 2.9. *The wealth corresponding to the trading speed $\theta \in \mathcal{T}$ is defined by*

$$W^\theta_t := \int_0^t P^\theta_s \theta_s ds.$$

On the other hand, the remaining inventory is defined by

$$Q^\theta_t := q_0 - \int_0^t \theta_s ds,$$

where $q_0 > 0$ is the initial inventory. We define the value function $\mathcal{V}^\theta$ by

$$\mathcal{V}^\theta := W_T^\theta - \phi \int_0^T (Q_t^\theta)^2 dt + Q_T^\theta (P_T^\theta - \alpha Q_T^\theta) \tag{2.4}$$

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

In this paper we study the optimal execution problem given by the optimization problem

$$\sup_{\theta \in \mathcal{T}} \mathbb{E}[\mathcal{V}^\theta]. \tag{2.5}$$

The first term of the value function indicates that the trader would like to maximize the wealth acquired by following the trading strategy θ. If the investor arrives at the terminal time with a nonzero inventory Q_T^θ, the third term of the value function $Q_T^\theta(P_T^\theta - \alpha Q_T^\theta)$ ensures that the remaining inventory is executed with a penalization $\alpha \geq 0$. Finally, the term $-\phi \int_0^T (Q_t^\theta)^2 dt$ penalizes 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 [11, 8]). In any case, a large value of the parameter 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 [9, 26, 11, 25, 18, 10, 19], but the authors would like to emphasize that the methodology proposed in this paper also applies to other alternative definitions of the value function, and we are not restricted to this particular choice of $\mathcal{V}^\theta$.

As it turns out, solving (2.5) over $\mathcal{T}_{sig}$ will be easier than doing so over the larger space $\mathcal{T}$. Due to the density result stated in Lemma B.3, we restrict the space of trading speeds from $\mathcal{T}$ to $\mathcal{T}_{sig}$, so that we solve the following problem instead:

$$\sup_{\theta \in \mathcal{T}_{sig}} \mathbb{E}[\mathcal{V}^\theta]. \tag{2.6}$$

3. Optimal execution.

The value function (2.4) is a nonlinear function of the underlying price process. However, for signature trading strategies $\theta \in \mathcal{T}_{sig}$ the cost function is a linear function on the signature of the midprice process. This is due to the shuffle product property (A.1)—each term in the value function can be rewritten as a linear function on the signature of the midprice process.

Lemma 3.1. *Let $\theta \in \mathcal{T}_{sig}$ be the signature trading speed given by $\theta_t = \langle \ell, \widehat{\mathbb{X}}_{0,t}^{<\infty} \rangle$ with ℓ a linear functional (Definition 2.4). Then, recalling the notation $\sqcup$ introduced in (2.2) and the notation 1^* for the constant linear functional introduced in Definition 2.4, we have*

1. $W_t^\ell = \left\langle \left[(e_{(2)} + 1^* - g^\ell) \sqcup\!\sqcup \ell \right] \sqcup (1), \widehat{\mathbb{X}}_{0,t}^{<\infty} \right\rangle$;
2. $Q_t^\ell = \langle q_0 1^* - \ell \sqcup (1), \widehat{\mathbb{X}}_{0,t}^{<\infty} \rangle$;
3. $\int_0^t (Q_s^\ell)^2 ds = \langle [q_0 1^* - \ell \sqcup (1)]^{\sqcup\!\sqcup 2} \sqcup (1), \widehat{\mathbb{X}}_{0,t}^{<\infty} \rangle$;
4. $Q_t^\ell (P_t^\ell - \alpha Q_t^\ell) = \langle [q_0 1^* - \ell \sqcup (1)] \sqcup\!\sqcup [e_{(2)} + 1^* - g^\ell] - \alpha [q_0 1^* - \ell \sqcup (1)]^{\sqcup\!\sqcup 2}, \widehat{\mathbb{X}}_{0,t}^{<\infty} \rangle$.

Therefore, the optimal liquidation problem (2.5) is transformed into the following problem.

Proposition 3.2. *Let $\theta \in \mathcal{T}_{sig}$ be the signature trading speed given by $\theta_t = \langle \ell, \widehat{\mathbb{X}}^{<\infty}_{0,t} \rangle$ with ℓ a linear functional (Definition 2.4). Then, the value function can be written as*

$$\mathcal{V}^\theta = \Big\langle \Big[(e_{(2)} + 1^* - g^\ell) ⧢ \ell \Big] \sqcup (1) - \phi [q_0 1^* - \ell \sqcup (1)]^{⧢ 2} \sqcup (1) - \alpha [q_0 1^* - \ell \sqcup (1)]^{⧢ 2} + [q_0 1^* - \ell \sqcup (1)] ⧢ (e_{(2)} + 1^* - g^\ell), \widehat{\mathbb{X}}^{<\infty}_{0,T} \Big\rangle .$$

Hence, the optimal liquidation problem (2.5) is reduced to

$$\sup_\ell \Big\langle \Big[(e_{(2)} + 1^* - g^\ell) ⧢ \ell \Big] \sqcup (1) - \phi [q_0 1^* - \ell \sqcup (1)]^{⧢ 2} \sqcup (1) - \alpha [q_0 1^* - \ell \sqcup (1)]^{⧢ 2} + [q_0 1^* - \ell \sqcup (1)] ⧢ (e_{(2)} + 1^* - g^\ell), \mathbb{E}\Big[\widehat{\mathbb{X}}^{<\infty}_{0,T} \Big] \Big\rangle . \tag{3.1}$$

The value function $\mathcal{V}^\theta$ depends on two aspects: a stochastic component and the control θ. Moreover, this dependency is nonlinear. Proposition 3.2 separates this dependency into a deterministic component that solely depends on the control, and 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 (2.4) is moved inside the linear functional—in other words, the resulting optimization problem (3.1) 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 [13] that under certain growth assumptions, the expected signature determines the law of the stochastic process. Therefore, the optimization problem depends on the entire law of the process because (3.1) depends on the expected signature of the midprice process.

3.1. Numerically solving the optimal execution problem. The optimization problem (3.1) in Proposition 3.2 involves the full expected signature $\mathbb{E}[\widehat{\mathbb{X}}^{<\infty}_{0,T}]$. In practice, however, one considers the truncated expected signature of order $N \in \mathbb{N}$, i.e., $\mathbb{E}[\widehat{\mathbb{X}}^{\leq N}_{0,T}]$.

However, the fast decay of the signature—$\widehat{\mathbb{X}}^{\mathbf{w}}_{0,T}$ decays factorially in the length of $\mathbf{w}$—implies that the first few terms dominate the rest, and not much information is lost in the truncation. As a consequence, the expected signature typically also decays factorially ([13] shows this for wide classes of Lévy, Markov, and Gaussian processes).

Figure 1 shows $\sqrt{\sum_{\mathbf{w} \in \mathcal{I}, |\mathbf{w}|=N} (\widehat{\mathbb{X}}^{\mathbf{w}}_{0,T})^2}$, i.e., the size of the Nth level of the signature, 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 it is sufficient to consider the first few signature terms only.

A rigorous analysis of the error arising from approximating a trading speed by a signature trading speed on the truncated signature would be of obvious practical use, but it is outside of the scope of the paper.

Once the signature is truncated at a certain level $N \in \mathbb{N}$, the optimization problem (3.1) consists of finding the global maximum of a certain polynomial in several variables. For

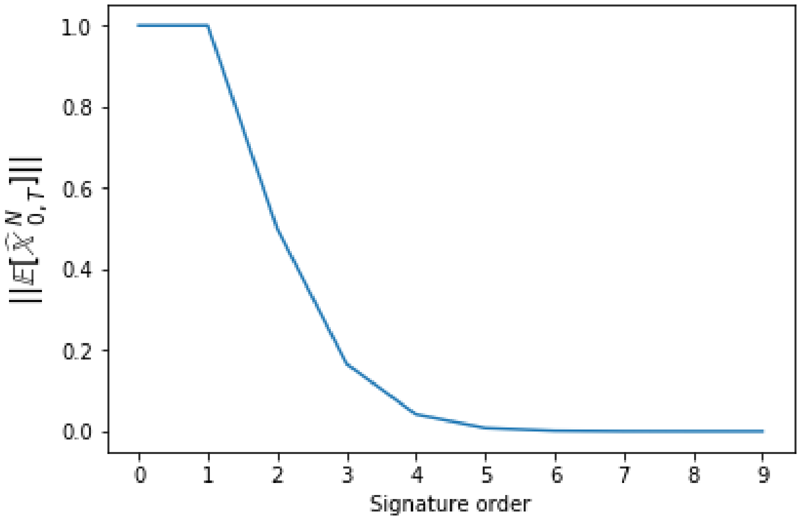

Figure 1. $\sqrt{\sum_{\mathbf{w}\in\mathcal{I},|\mathbf{w}|=N}(\widehat{\mathbb{X}}^{\mathbf{w}}_{0,T})^2}$ *as a function of the signature order N in the case where the midprice process is a Brownian motion. The factorial decay of the signature makes higher order terms small compared to the first few terms.*

example, one can show that if a linear permanent and temporary market impact is considered, then the polynomial is quadratic, and finding the optimal trading speed reduces to finding the (unique) global maximum of a quadratic polynomial in several variables.

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 realization can be computed using publicly available software such as `esig`[1] or `iisignature`.[2]

4. Extensions.

In section 3, we studied a specific optimal liquidation problem. Here, we analyze different extensions to the liquidation problem, and we study how they fit in our framework.

4.1. Modelling the execution price with exogeneous information.

For $\theta \in \mathcal{T}_{sig}$, in section 2.4 the market impact was defined as a function of the trading speed and the unaffected midprocess:

$$P^\theta_t := X_t - \langle g^\theta, \widehat{\mathbb{X}}^{<\infty}_{0,t}\rangle. \tag{4.1}$$

However, there are other factors that affect the impact of a trading order [30, 34]. For instance, one may want to incorporate the total traded volume $V : [0,T] \to \mathbb{R}$ in the market impact [34]. 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 [30, 35, 29].

[1] https://pypi.org/project/esig/.
[2] https://pypi.org/project/iisignature/.

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

$$P_t^\theta := \langle f^\theta, \widehat{\mathbb{Z}}_{0,t}^{<\infty}\rangle, \tag{4.2}$$

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 $Y_t^1, \dots, Y_t^n$ the midprice processes of n alternative assets that the trader believes affect the execution price of the main asset. Notice that (4.1) is a particular case of (4.2). Other exogenous information may also be added to $\widehat{Z}$. In this case, the signature trading strategy will depend on the signature $\widehat{\mathbb{Z}}_{0,t}^{<\infty}$, allowing the trader to use the exogeneous information.

The methodology proposed in this paper still applies to this setting: the optimization problem (2.5), for the new definition of market impact, is reduced to an optimization problem similar to (3.1), namely,

$$\begin{aligned}\sup_{\ell} \Big\langle \Big(f^\ell \sqcup\!\sqcup \ell\Big) \sqcup (1) - \phi[q_0 1^* - \ell \sqcup (1)]^{\sqcup\!\sqcup 2} \sqcup (1) - \alpha[q_0 1^* - \ell \sqcup (1)]^{\sqcup\!\sqcup 2} \\ + [q_0 1^* - \ell \sqcup (1)] \sqcup\!\sqcup f^\ell, \mathbb{E}\Big[\widehat{\mathbb{Z}}_{0,T}^{<\infty}\Big]\Big\rangle.\end{aligned} \tag{4.3}$$

4.2. Optimal trading, as opposed to liquidation. In this paper 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 paper's framework can be modified for this purpose by setting $q_0 = 0$ in Definition 2.9.

4.3. Cross-asset portfolio liquidation. The discussion in section 4.1 suggests another extension of the original problem studied in this paper. Suppose there are n assets $Y^1, \dots, Y^n$ and a trader has an initial portfolio $q = (q_1, \dots, q_n) \in \mathbb{R}_+^n$. If the trader wishes to liquidate the inventory q (see [35, 29]), she can consider an optimal control problem similar to (4.3) 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.

4.4. Other value functions. The value function considered in (2.5) was chosen to be consistent with the literature [9, 26, 11, 25, 18, 10, 19]. 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. 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 [12, section 6.5], [26], [11] 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 realizations of the process. Once the expected signature is estimated, we solve (3.1) to obtain a signature trading strategy—what parameters have been used in the optimization problem will be specified on each numerical experiment. Finally, we evaluate the performance of the constructed signature trading strategy by sampling 10,000 new realizations 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 [12, 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 optimization problem (3.1) that were used. However, the performance of the signature method does not appear to be sensitive with respect to the choice of these parameters in the sense that different choice of parameters also gives similarly satisfactory results.

5.1. Brownian motion with temporary and permanent market impact. In this section we consider the framework studied in [12, section 6.5]. 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 $\theta \in \mathcal{T}_{sig}$, the execution price is given by a *permanent* market impact and a *temporary* market impact:

$$P_t^\theta := X_t - k \int_0^t \theta_s ds - \lambda \theta_t$$

with $k, \lambda > 0$.

Next, we solve (3.1) with truncated signatures of order 7. The parameters are $q_0 = 1$, $\lambda = 10^{-3}$, $k = 10^{-4}$, $\alpha = 10$, $\sigma = 0.02$, and $T = 1$, and different values for ϕ are considered. As shown in the literature (see [12, section 6.5]) the optimal trading speed does not depend on the midprice. Moreover, if we set $\phi = 0$ there is no running inventory penalty, and it is known that the optimal trading speed is constant. On the other hand, as the value of ϕ increases, the trader liquidates the inventory faster. All these features are captured in the results we obtained—see Figure 2.

5.2. Incorporating order-flow. In [11], 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 market buying and selling orders of all market participants, respectively, i.e., the difference $\mu^+ - \mu^-$ is the net order flow. We assume the market orders follow the dynamics

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

with $L_t^\pm$ independent Poisson processes of intensity λ_0, and $\eta_i^\pm \sim Exp(\eta_0 \kappa)$ having an exponential distribution. Moreover, a temporary market impact was included as well—i.e., the execution price is given by (2.3) with $g^\ell := \lambda \ell$, $\lambda > 0$.

Figure 3 shows the inventory for 100 realizations of the midprice path, both for the signature trading speed and the optimal trading speed that was derived in [11]. Table 1 shows

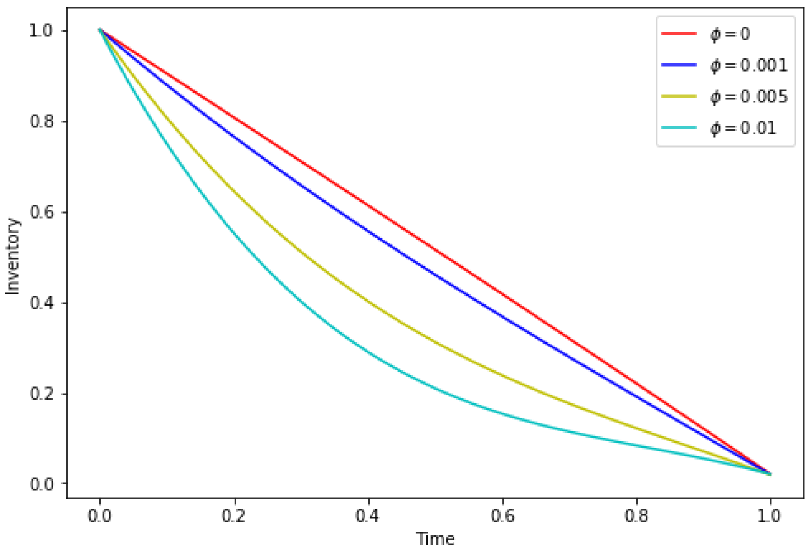

Figure 2. *The trader's inventory for* 100 *midprice path realizations and the setting considered in section* 5.1*. Different running inventory penalties* ϕ *were considered.*

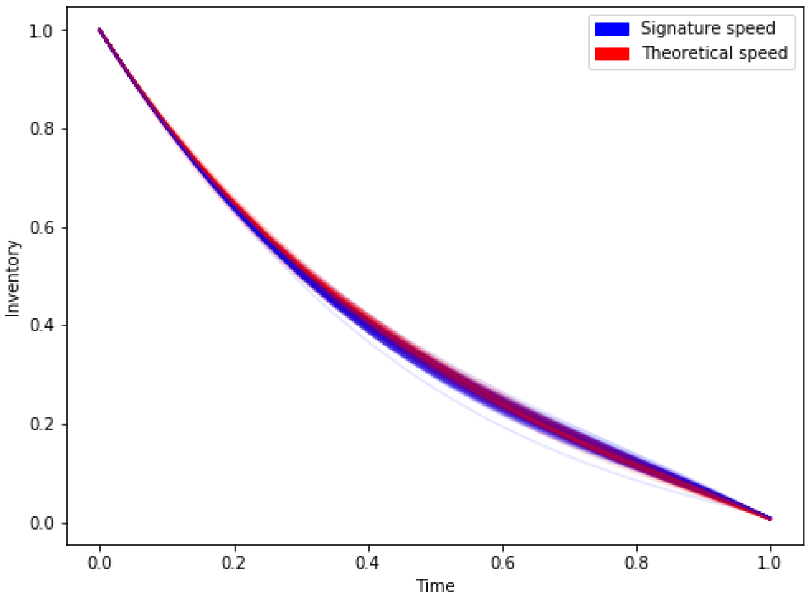

Figure 3. *The trader's inventory for* 100 *midprice path realizations and the setting considered in section* 5.2*, both for the theoretical optimal speed (red) and the signature trading speed (blue).*

the performance of the signature trading strategy compared to the optimal trading speed and the TWAP strategy. The parameters we considered are $\lambda = 5 \cdot 10^{-4}$, $k = 10^{-4}$, $q_0 = 1$, $\alpha = 2$, $\phi = 5 \cdot 10^{-3}$, $\sigma = 0.1$, $\kappa = \lambda_0 = 5$, $\eta_0 = 0.8$, and signatures of order 7.

Table 1

Performance of the theoretical optimal speed, signature trading speed, and the TWAP for the setting considered in section 5.2. We use the expected value function, estimated using the testing set, to measure performance.

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

5.3. Incorporating trading signals. Lehalle and Neuman in [26] 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.

In this case, the midprice process is $X_t := \int_0^t I_s ds + \sigma W_t$, where I is the signal process, $\sigma > 0$ is volatility, and W is a Brownian motion. In the original paper [26], the signal I follows an Ornstein–Uhlenbeck process, i.e., $dI_t = -\gamma I_t dt + \sigma_0 dW_t$, where $\gamma, \sigma > 0$ are constants. Therefore, given that the midprice process is a semimartingale, this example also falls within our framework. Note that this setting is a particular case of the one studied in [11].

[26] considers a linear temporary price impact; thus, the execution price is given by (2.3), where $g^\ell := \lambda\ell$ with $\lambda > 0$.

Figure 4 shows the running inventory for 100 realizations of the midprice process, both for the signature trading speed and the optimal trading speed that was derived in [26]. The parameters are $q_0 = 1$, $\lambda = 10^{-3}$, $\alpha = 10^{-2}, \phi = 10^{-3}, I_0 = 0.02$, and $\gamma = 0.1$. Truncated

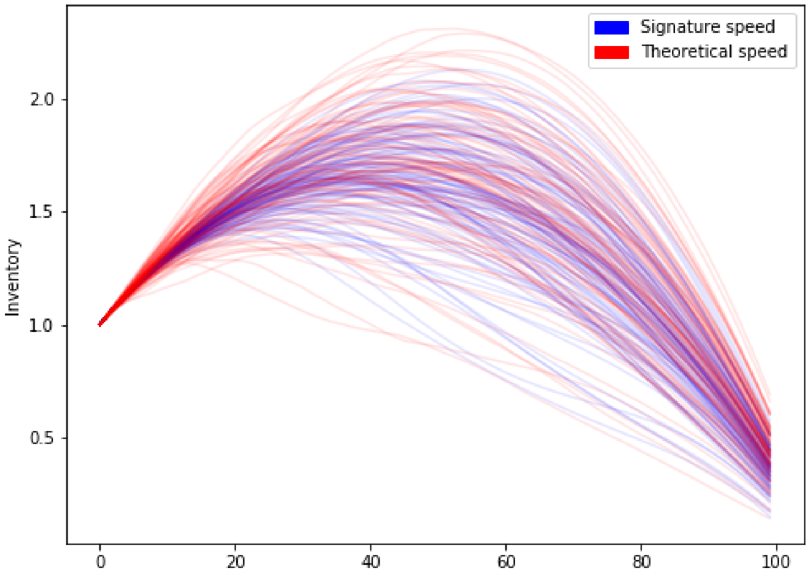

Figure 4. *The trader's inventory for* 100 *midprice path realizations and the setting considered in section* 5.3*, both for the theoretical optimal speed (red) and the signature trading speed (blue).*

Table 2

Performance of the theoretical optimal speed, signature trading speed, and the TWAP for the setting considered in section 5.3. We use the expected value function, estimated using the testing set, to measure performance.

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

signatures of order 9 are considered. As we see, the signature trading speed seems to be a good approximation to the theoretical optimal speed. This is also reflect in the performance of the signature trading strategy with respected to the theoretical optimal speed—see Table 2.

Notice that the presence of the signal in the midprice process introduces a positive drift, and therefore, as illustrated by Figure 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.

5.4. 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

$$P_t^\theta := \sigma W_t^H - \lambda\theta_t,$$

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

Figure 5 shows the midprice and inventory when $H = 1/3$, $\sigma = 0.02$, $q_0 = 1$, $\phi = 0$, $\lambda = 10^{-3}$, $\alpha = 0.1$, $T = 1$, and truncated signatures of order 7 are considered.

As we see, the behavior 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 illustrates that this is not the case for the fractional Brownian motion, and the trading speed

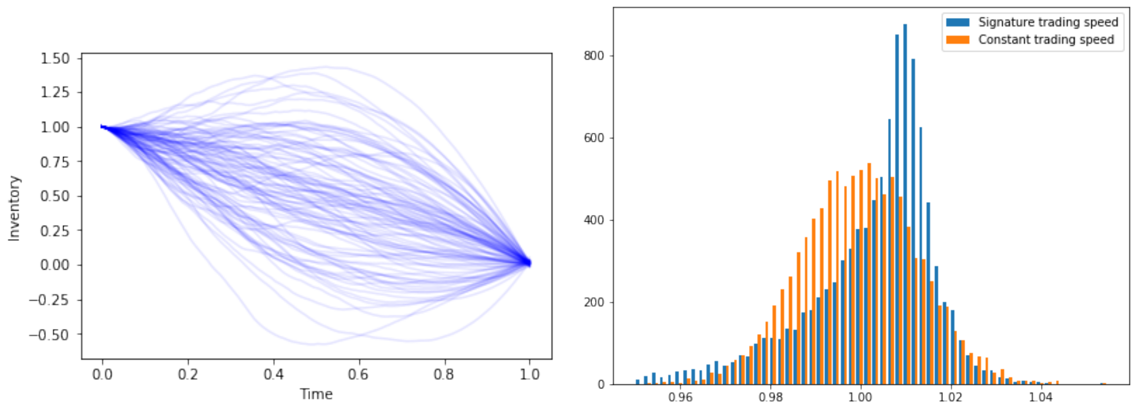

Figure 5. *Trader's inventory (left) and the trader's wealth distribution (right) in the case where the midprice process is a fractional Brownian motion with $H = 1/3$.*

Table 3

Performance of the signature trading speed and the TWAP for the fractional Brownian motion, for two choices of Hurst parameter H. We use the expected value function, estimated using the testing set, to measure performance.

Hurst parameter H	**Signature trading speed**	**TWAP**
1/3	$1.0031498 \pm 1.38 \times 10^{-4}$	$0.9991785 \pm 5.73 \times 10^{-4}$
0.7	$1.0203925 \pm 2.82 \times 10^{-6}$	$0.99921570 \pm 1.08 \times 10^{-6}$

depends strongly on the midprice process. In fact, if we look at the expected value function 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 the TWAP (see Table 3 for $H = 1/3$ and $H = 0.7$). This outperformance of the signature trading speed is reflected in the wealth distribution of both strategies shown in Figure 5 (right). This is not surprising, as the fractional Brownian motion with $H \neq 1/2$ exhibits a memory effect and an optimal trading speed should be able to exploit that.

6. Experiments with market data.

To solve (3.1), the only information we need about the midprice process is its expected signature. In this section, we use real market data to estimate the expected signature, which we use to solve (3.1). Then, we evaluate the performance of the optimal execution strategy in an out-of-sample set of price paths.

We consider 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.[3] This data was divided into a *training set* of 10 months (January–October) and a *testing set* of 2 months (November–December).

We consider 15-minute windows from different times of each trading day—more specifically, we consider 10:00–10:15, 11:00–11:15, 12:00–12:15, and 13:00–13:15. We estimate 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 training set. Therefore, to some extent, we assume that the midprice process follows a similar behavior 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 optimization problem (3.1) to estimate the optimal signature trading speed. We include a temporary and market impact:

$$P_t^\theta := X_t - k \int_0^t \theta_s ds - \lambda \theta_t.$$

Other parameters are $\lambda = 10^{-3}$, $k = 10^{-4}$, $\alpha = 0.1, \phi = 10^{-4}$, and $q_0 = 1$. We evaluate the performance on the testing set for each of the 15-minute windows. Following [11], we compare the performance against the Almgren–Chriss execution strategy [1]. More specifically, we

[3] https://lobsterdata.com/.

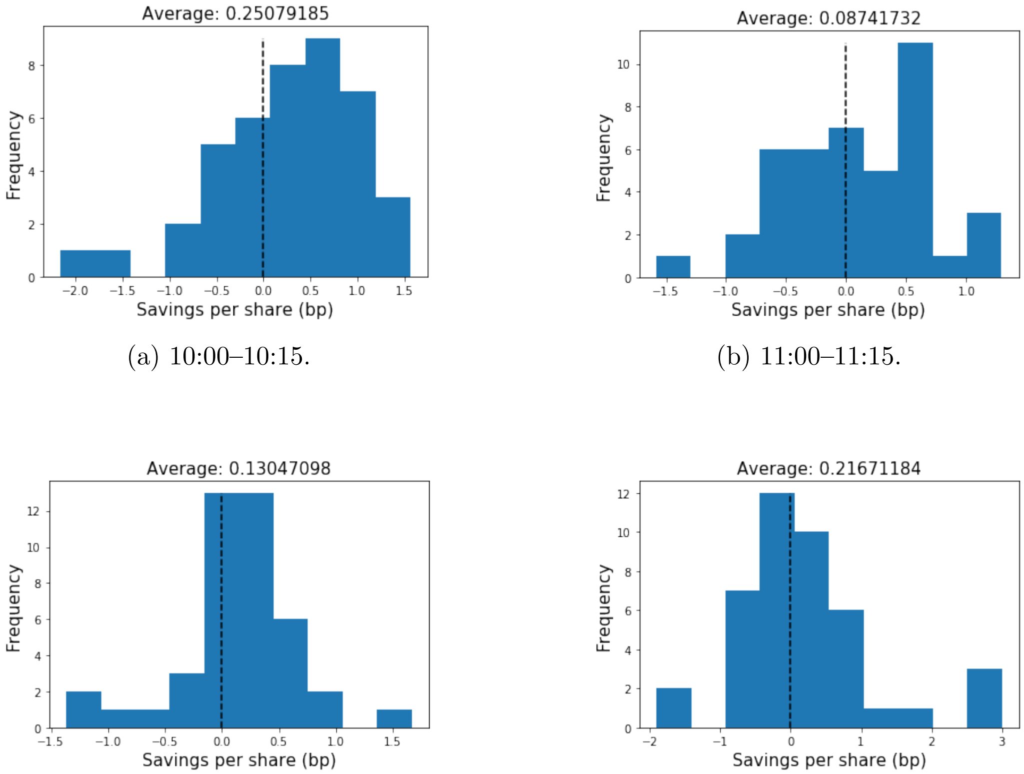

(a) 10:00–10:15.

(b) 11:00–11:15.

(c) 12:00–12:15.

(d) 13:00–13:15.

Figure 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.*

consider the *savings per share* metric (in basis points) used in [11], 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.

The results are shown in Figure 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-agnostic: we can, in a nonparametric and model-agnostic way, estimate the optimal trading speed from market data.

7. Conclusion.

In this paper we proposed 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 such as [26] and [11].

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

We showed that in the cases where the optimal trading speed is known, the signature-based numerical approach delivered similar results. Moreover, the generality of the approach allows the estimation of the optimal trading speed in those settings where the optimal solution is unknown. We also showed how our methodology can be employed with real market data, and we demonstrated that the signature approach outperforms the Almgren–Chriss benchmark.

Appendix A. Rough paths preliminaries.

The aim of this section is to introduce all the aspects of rough paths theory that are used in the article. For a more detailed introduction to the theory of rough paths, the authors refer the reader to [28, 23].

A.1. Tensor algebra.

A rough path is a path that takes value on a certain graded space, called the tensor algebra. This subsection will introduce these algebras, as well as another crucial space—the dual space of the tensor algebra.

Definition A.1 (extended tensor algebra). *Let $d \geq 1$. We denote by $T((\mathbb{R}^d))$ the extended tensor algebra over $\mathbb{R}^d$, which is defined by*

$$T((\mathbb{R}^d)) := \{\mathbf{a} = (a_0, a_1, \ldots, a_n, \ldots) \mid a_n \in (\mathbb{R}^d)^{\otimes n}\},$$

where $\otimes$ denotes the tensor product. Given $\mathbf{a} = (a_i)_{i=0}^{\infty}, \mathbf{b} = (b_i)_{i=0}^{\infty} \in T((\mathbb{R}^d))$, define the sum $+$ and product $\otimes$ by

$$\mathbf{a} + \mathbf{b} := (a_i + b_i)_{i=0}^{\infty},$$
$$\mathbf{a} \otimes \mathbf{b} := \left(\sum_{k=0}^{i} a_k \otimes b_{i-k} \right)_{i=0}^{\infty}.$$

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

Similarly, we can define the tensor algebra and truncated tensor algebra as the space of all finite sequences and all sequences of a given length, respectively.

Definition A.2. *The tensor algebra over $\mathbb{R}^d$, denoted by $T(\mathbb{R}^d) \subset T((\mathbb{R}^d))$, is given by*

$$T(\mathbb{R}^d) := \{\mathbf{a} = (a_i)_{i=0}^{\infty} \mid a_i \in (\mathbb{R}^d)^{\otimes i} \text{ and } \exists N \in \mathbb{N} \text{ such that } a_i = 0 \, \forall n \geq N\}.$$

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

$$T^{(N)}(\mathbb{R}^d) := \{\mathbf{a} = (a_i)_{i=0}^{\infty} \mid a_i \in (\mathbb{R}^d)^{\otimes i} \text{ and } a_i = 0 \, \forall i \geq N\}.$$

Let $\{e_1, \ldots, e_d\} \subset \mathbb{R}^d$ be a basis for $\mathbb{R}^d$. This induces a dual basis $\{e_1^*, \ldots, e_d^*\} \subset (\mathbb{R}^d)^*$ for $(\mathbb{R}^d)^*$, where $(\mathbb{R}^d)^*$ denotes the dual space of $\mathbb{R}^d$—i.e., the space of all continuous linear functions $\mathbb{R}^d \to \mathbb{R}$. We may define a basis for $(\mathbb{R}^d)^{\otimes n}$ by

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

Similarly, a basis of $((\mathbb{R}^d)^*)^{\otimes n}$ is defined by

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

This induces, in a natural way, a basis for $T((\mathbb{R}^d))$ and $T((\mathbb{R}^d)^*)$.

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

Example A.3. Consider the following examples for $\mathbb{R}^2$.
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 $\mathbf{w} = \mathbf{i_1} \ldots \mathbf{i_n}, \mathbf{v} = \mathbf{j_1} \ldots \mathbf{j_m} \in \mathcal{W}(\mathcal{A}_d)$ be two words. Their sum is the formal sum $\mathbf{w} + \mathbf{v} \in \mathcal{W}(\mathcal{A}_d)$. Their concatenation, on the other hand, is defined by

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

These two operations induce analogous operations on $T((\mathbb{R}^d)^*)$, and with some abuse of notation we will even use concatenation on $\mathcal{W}(\mathcal{A}_d)$ and $T((\mathbb{R}^d)^*)$ interchangeably—i.e., we will sometimes write $\ell \mathbf{w} \in T((\mathbb{R}^d)^*)$ for $\ell \in T((\mathbb{R}^d)^*)$ 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 $\mathbf{w}$.

Example A.4. 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 paper: the *shuffle product* $\sqcup\!\sqcup$. Intuitively, the shuffle product accounts for all the possible ways of riffle shuffling two decks of cards. The precise definition is given below.

Definition A.5 (shuffle product). *The shuffle product* $\sqcup\!\sqcup : \mathcal{W}(\mathcal{A}_d) \times \mathcal{W}(\mathcal{A}_d) \to \mathcal{W}(\mathcal{A}_d)$ *is defined inductively by*

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

$$\mathbf{w} \sqcup\!\sqcup \varnothing = \varnothing \sqcup\!\sqcup \mathbf{w} = \mathbf{w}$$

for all words $\mathbf{u}, \mathbf{v}$ *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* $T((\mathbb{R}^d)^*)$ *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 $f, g \in T((\mathbb{R}^d)^*)$.

Example A.6. We have the following.

1. $\mathbf{12} \sqcup\!\sqcup \mathbf{3} = \mathbf{123} + \mathbf{132} + \mathbf{312}$.
2. $\mathbf{12} \sqcup\!\sqcup \mathbf{23} = 2 \cdot \mathbf{1223} + \mathbf{1232} + \mathbf{2123} + \mathbf{2132} + \mathbf{2312}$.

Definition A.7. *Let $Q \in \mathbb{R}[x]$ be a polynomial on one variable. Write $Q(x) = a_0 + a_1 x + a_2 x^2 + \cdots + a_n x^n$. Then, Q induces the map $Q^{\sqcup\!\sqcup} : T((\mathbb{R}^d)^*) \to T((\mathbb{R}^d)^*)$ given by*

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

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

A.2. Rough paths. We will now define a crucial object in this paper: the signature of a path.

Definition A.8 (signature of a path). *Let $0 \le s < t \le T$. For a piecewise smooth path $X : [0,T] \to \mathbb{R}^d$, we define the signature of X over $[s,t]$ by*

$$\mathbb{X}^{<\infty}_{s,t} := (1, \mathbb{X}^1_{s,t}, \ldots, \mathbb{X}^n_{s,t}, \ldots) \in T((\mathbb{R}^d)),$$

where

$$\mathbb{X}^n_{s,t} := \int_{s<u_1<\ldots<u_n<t} dX_{u_1} \otimes \cdots \otimes dX_{u_n} \in (\mathbb{R}^d)^{\otimes n}.$$

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

$$\mathbb{X}^{\le N}_{s,t} := (1, \mathbb{X}^1_{s,t}, \ldots, \mathbb{X}^N_{s,t}) \in T^{(N)}(\mathbb{R}^d).$$

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

Example A.9. Throughout this paper, we will constantly work with linear functions on the signature. Therefore, it will be useful to see a few examples that will be used in other sections.

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

1. $\langle \mathbf{2}, \mathbb{X}^{<\infty}_{0,T} \rangle = \int_0^T dX^2_t = X^2_T - X^2_0$.
2. $\langle \varnothing, \mathbb{X}^{<\infty}_{0,T} \rangle = 1$.
3. $\langle \mathbf{21}, \mathbb{X}^{<\infty}_{0,T} \rangle = \int_0^T \int_0^t dX^2_s dX^1_t = \int_0^T (X^2_t - X^2_0) dX^1_t$.
4. Let $\ell \in T((\mathbb{R}^2)^*)$. Then, $\langle \ell \mathbf{1}, \mathbb{X}^{<\infty}_{0,T} \rangle = \int_0^T \langle \ell, \mathbb{X}^{<\infty}_{0,t} \rangle dX^1_t$.

Definition A.10 (geometric p-rough paths). *Let $T > 0$ and $p \ge 1$. Denote by $\lfloor p \rfloor$ the integer part of $p \in \mathbb{R}$. Let $\Delta_T := \{(s,t) \in [0,T] \times [0,T] \mid s \le t\}$. A function $\mathbb{X} : \Delta_T \to T^{(\lfloor p \rfloor)}(\mathbb{R}^d)$ is said to be a geometric p-rough path if it is the limit (under the p-variation distance* [28, Definition 1.5]*) of signatures of order $\lfloor p \rfloor$ of piecewise smooth paths. The space of all geometric p-rough paths will be denoted by $G\Omega_p([0,T]; \mathbb{R}^d)$.*

Each $\mathbb{X} = (1, \mathbb{X}^1, \ldots, \mathbb{X}^{\lfloor p \rfloor}) \in G\Omega_p([0,T]; \mathbb{R}^d)$ can be uniquely extended to a N-geometric rough path for any $N \ge p$ [28, Theorem 3.7]. Analogously to the smooth case, the full extension $\mathbb{X}^{<\infty} = (1, \mathbb{X}^1, \ldots, \mathbb{X}^N, \ldots)$ will be defined as the *signature* of $\mathbb{X}$.

Many stochastic processes that are used in the literature [12, section 6.5], [26], [11] are almost surely geometric rough paths. For example, the signature of a semimartingale, defined using Stratonovich integration, is almost surely a geometric p-rough path for any $p \in (2,3)$ [27, section 2.4.2]. 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 $p > 1/H$ [17]. We will now state some properties of signatures that will be useful in this article.

Lemma A.11 (shuffle product property [28, Theorem 2.15]). *Let $\mathbb{X} \in G\Omega_p([0,T];\mathbb{R}^d)$ be a geometric p-rough path. Let $\ell_1, \ell_2 \in T((\mathbb{R}^d)^*)$ be two linear functionals. Then,*

$$\langle \ell_1, \mathbb{X}^{<\infty}_{0,T}\rangle\langle \ell_2, \mathbb{X}^{<\infty}_{0,T}\rangle = \langle \ell_1 \sqcup\!\sqcup \ell_2, \mathbb{X}^{<\infty}_{0,T}\rangle \quad \forall \ell_1, \ell_2 \in T((\mathbb{R}^d)^*). \tag{A.1}$$

The shuffle product will be extensively used throughout this paper. It guarantees that the product of two linear functions on the signature is another linear function on the signature, which is given explicitly in terms of the shuffle product.

Remark A.12. It turns out [6] that the signature $\mathbb{X}^{<\infty}_{0,T}$ completely characterizes $\mathbb{X}$—up to the so-called tree-like equivalences (see [6, Definition 1.1]). Moreover, we have the following corollary (see [14, Corollary 5.7]).

Corollary A.13. *Let $\mathbb{X} \in G\Omega_p([0,T];\mathbb{R}^d)$. If there exists a projection of $\mathbb{X}$ that is strictly monotone, then the signature $\mathbb{X}^{<\infty}_{0,T}$ determines $\mathbb{X}$ up to translations.*

Appendix B. Trading speeds and signature trading speeds.

This section will introduce a more formal treatment of trading strategies and signature trading strategies, introduced in section 2.3, and will conclude with the proof of Lemma B.3.

B.1. Notation.

Given a continuous path $X \in C([0,T];\mathbb{R})$, we will denote its augmentation by $\widehat{X} \in C([0,T];\mathbb{R}^2)$ defined by $\widehat{X}_t := (t, X_t) \in \mathbb{R}^2)$. Let $p \geq 1$. Recalling the notion of *geometric rough path* introduced in Appendix A, we define for $t \in [0,T]$

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

where the closure is taken under d_{p-var}, i.e., the p-variation distance (see [28, Definition 1.5]). Given $\widehat{\mathbb{X}} \in \widehat{\Omega}^p_t$, we will write by $X \in C([0,T];\mathbb{R})$ the unaugmented coordinate process.

Intuitively, elements of $\widehat{\Omega}^p_t$ are signatures of paths of the form (u, X_u), 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 A.13 that $\widehat{\mathbb{X}}^{<\infty}_{0,t}$ completely characterizes $\widehat{\mathbb{X}} \in \widehat{\Omega}^p_t$ (and hence X).

The space $\widehat{\Omega}^p_T$ will be our space of price paths. We will equip it with a probability space $(\widehat{\Omega}^p_T, \mathcal{B}(\widehat{\Omega}^p_T), \mathbb{P})$. Given a rough path $\widehat{\mathbb{X}} \in \widehat{\Omega}^p_T$, 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. This is equivalent to the framework discussed in section 2.3.

B.2. Trading speeds.

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

Definition B.1 (trading speeds). *Define the metrizable space $\Lambda_T := \bigcup_{t\in[0,T]} \widehat{\Omega}^p_t$. 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 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 (nonanticipative) 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 [24, 15, 2, 20, 4, 32], and a similar definition of trading strategies was considered in [32].

We define the class of signature trading speeds, on the other hand, as follows.

Definition B.2 (signature trading speeds). *The space of signature trading speeds $\mathcal{T}_{sig} \subset \mathcal{T}$ is defined by*

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

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

B.3. Approximation of trading speeds by signature trading speeds. 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:

Lemma B.3. *Let $\varepsilon > 0$. Then, there exists a compact set $\mathcal{K} \subset \widehat{\Omega}^p_T$ such that*

1. $\mathbb{P}[\mathcal{K}] > 1 - \varepsilon$;
2. $\mathcal{T}_{sig}$, *restricted to* $\mathrm{Hist}(\mathcal{K}) \subset \Lambda_T$, *is dense in* $\mathcal{T}$, *where*

$$\mathrm{Hist}(\mathcal{K}) := \left\{\widehat{\mathbb{X}}|_{[0,t]} \in \widehat{\Omega}_t \mid \widehat{\mathbb{X}} \in \mathcal{K}, t \in [0,T]\right\} \subset \Lambda_T$$

is the history of all paths (i.e., all stopped paths) in $\mathcal{K}$.

Proof. Let $\varepsilon > 0$. Because $\widehat{\Omega}^p_T$ is a Polish space [22, Theorem 23], there exists $\mathcal{K} \subset \widehat{\Omega}^p_T$ compact such that $\mathbb{P}[\mathcal{K}] > 1 - \varepsilon$ [7, Theorem 7.1.7].

Let $\theta_1, \theta_2 \in \mathcal{T}_{sig}$. Then, by definition there exist linear functionals $\ell_1, \ell_2 \in T((\mathbb{R}^2)^*)$ such that $\theta_i(\widehat{\mathbb{X}}|_{[0,t]}) = \langle \ell_i, \widehat{\mathbb{X}}^{<\infty}_{0,t}\rangle$ for all $\widehat{\mathbb{X}}|_{[0,t]} \in \Lambda_T, i = 1,2$. Define $\theta(\widehat{\mathbb{X}}|_{[0,t]}) := \langle \ell_1 \sqcup\!\sqcup \ell_2, \widehat{\mathbb{X}}^{<\infty}_{0,t}\rangle$. Then, by the shuffle product property (A.1) we have

$$\begin{aligned}\theta_1(\widehat{\mathbb{X}}|_{[0,t]})\theta_2(\widehat{\mathbb{X}}|_{[0,t]}) &= \langle \ell_1, \widehat{\mathbb{X}}^{<\infty}_{0,t}\rangle\langle \ell_2, \widehat{\mathbb{X}}^{<\infty}_{0,t}\rangle \\ &= \langle \ell_1 \sqcup\!\sqcup \ell_2, \widehat{\mathbb{X}}^{<\infty}_{0,t}\rangle \\ &= \theta(\widehat{\mathbb{X}}|_{[0,t]}).\end{aligned}$$

Therefore, and because the sum of two signature trading speeds is trivially a signature trading speed, $\mathcal{T}_{sig}$ form an algebra. On the other hand, the uniqueness of the signature (Corollary A.13) implies that $\mathcal{T}_{sig}$ separates points. Indeed, given $\widehat{\mathbb{X}}|_{[0,t]}, \widehat{\mathbb{Y}}|_{[0,t]} \in \widehat{\Omega}^p_T$ distinct, because we have $\widehat{\mathbb{X}}^{<\infty}_{0,t} \neq \widehat{\mathbb{Y}}^{<\infty}_{0,t}$ we immediately have that there exists $\ell \in T((\mathbb{R}^2)^*)$ such that $\langle \ell, \widehat{\mathbb{X}}^{<\infty}_{0,t}\rangle \neq \langle \ell, \widehat{\mathbb{Y}}^{<\infty}_{0,t}\rangle$. Moreover, $\mathcal{T}_{sig}$ contains constants, as $\langle \varnothing, \widehat{\mathbb{X}}^{<\infty}_{0,t}\rangle = 1$ for all $\widehat{\mathbb{X}}|_{[0,t]} \in \widehat{\Omega}^p_T$. Therefore, by the Stone–Weierstrass theorem [33] we conclude that $\mathcal{T}_{sig}$, restricted to $\mathrm{Hist}(\mathcal{K})$, is dense in $\mathcal{T}$. ∎

Appendix C. Transient market impact.

In Example 2.8, we included the *transient market impact* as an example of a market impact included in our framework. Recall that for a signature trading speed $\theta \in \mathcal{T}_{sig}$ given by $\theta_t = \langle \ell, \widehat{\mathbb{X}}^{<\infty}_{0,t} \rangle$, we model the market impact by $\langle g^\theta, \widehat{\mathbb{X}}^{<\infty}_{0,t} \rangle$. The transient market impact, on the other hand, is given by $\int_0^t \exp(-\rho(t-s))\theta_s ds$. In this section, we'll see how such a market impact can be approximated by market impacts included in our framework.

First, notice that for each $u \in [0,T]$ and $n \in \mathbb{N}$, if we define for each $n \in \mathbb{N}$ the multi-index $\mathbf{w_n} := \underbrace{(1,\dots,1)}_{n}$, we have

$$\frac{u^n}{n!} = \langle e_{\mathbf{w}_n}, \widehat{\mathbb{X}}^{<\infty}_{0,u} \rangle.$$

Hence, by Taylor expanding the function $x \mapsto \exp(x)$ we have, for $N \in \mathbb{N}$ and $t \in [0,T]$,

$$\begin{aligned}
\int_0^t \exp(-\rho(t-s))\theta_s ds &= \exp(-\rho t)\int_0^t \exp(\rho s)\langle \ell, \widehat{\mathbb{X}}^{<\infty}_{0,s}\rangle ds \\
&\approx \left(\sum_{n=0}^N \frac{(-\rho)^n t^n}{n!}\right)\int_0^t \left(\sum_{n=0}^N \frac{\rho^n s^n}{n!}\right)\langle \ell, \widehat{\mathbb{X}}^{<\infty}_{0,s}\rangle ds \\
&= \left\langle \sum_{n=0}^N (-\rho)^n e_{\mathbf{w_n}}, \widehat{\mathbb{X}}^{<\infty}_{0,t}\right\rangle \int_0^t \left\langle \sum_{n=0}^N \rho^n e_{\mathbf{w_n}}, \widehat{\mathbb{X}}^{<\infty}_{0,s}\right\rangle \langle \ell, \widehat{\mathbb{X}}^{<\infty}_{0,s}\rangle ds \\
&= \left\langle \sum_{n=0}^N (-\rho)^n e_{\mathbf{w_n}}, \widehat{\mathbb{X}}^{<\infty}_{0,t}\right\rangle \int_0^t \left\langle \left(\sum_{n=0}^N \rho^n e_{\mathbf{w_n}}\right) ⧢ \ell, \widehat{\mathbb{X}}^{<\infty}_{0,s}\right\rangle ds \\
&= \left\langle \sum_{n=0}^N (-\rho)^n e_{\mathbf{w_n}}, \widehat{\mathbb{X}}^{<\infty}_{0,t}\right\rangle \left\langle \left[\left(\sum_{n=0}^N \rho^n e_{\mathbf{w_n}}\right) ⧢ \ell\right] \sqcup (1), \widehat{\mathbb{X}}^{<\infty}_{0,t}\right\rangle \\
&= \left\langle \left(\sum_{n=0}^N (-\rho)^n e_{\mathbf{w_n}}\right) ⧢ \left(\left[\left(\sum_{n=0}^N \rho^n e_{\mathbf{w_n}}\right) ⧢ \ell\right] \sqcup (1)\right), \widehat{\mathbb{X}}^{<\infty}_{0,t}\right\rangle \\
&=: \langle g^{\theta,N}, \widehat{\mathbb{X}}^{<\infty}_{0,t}\rangle.
\end{aligned}$$

In other words, we may approximate the transient impact with a market impact $\langle g^{\theta,N}, \widehat{\mathbb{X}}^{<\infty}_{0,t}\rangle$ included in the framework introduced in section 2.4.

Appendix D. Proof of Lemma 3.1.

Proof of Lemma 3.1. Let $t \in [0,T]$.

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

$$\begin{aligned}
W^\ell_t &= \int_0^t P^\ell_s \langle \ell, \widehat{\mathbb{X}}^{<\infty}_{0,s}\rangle ds = \int_0^t (X_s - \langle g^\ell, \widehat{\mathbb{X}}^{<\infty}_{0,s}\rangle)\langle \ell, \widehat{\mathbb{X}}^{<\infty}_{0,s}\rangle ds \\
&= \int_0^t \left\langle (e_{(2)} + 1^* - g^\ell) ⧢ \ell, \widehat{\mathbb{X}}^{<\infty}_{0,s}\right\rangle ds = \left\langle \left[(e_{(2)} + 1^* - g^\ell) ⧢ \ell\right] \sqcup (1), \widehat{\mathbb{X}}^{<\infty}_{0,t}\right\rangle.
\end{aligned}$$

2. This part follows from the fact that $\int_0^t \langle \ell, \widehat{\mathbb{X}}^{<\infty}_{0,s}\rangle ds = \langle \ell \sqcup (1), \widehat{\mathbb{X}}^{<\infty}_{0,t}\rangle$.

3. Using 2,

$$\int_0^t (Q_s^\ell)^2 ds = \int_0^t \left\langle [q_0 1^* - \ell \sqcup (1)]^{\sqcup\!\sqcup 2}, \widehat{\mathbb{X}}_{0,s}^{<\infty} \right\rangle ds = \left\langle [q_0 1^* - \ell \sqcup (1)]^{\sqcup\!\sqcup 2} \sqcup (1), \widehat{\mathbb{X}}_{0,t}^{<\infty} \right\rangle .$$

4. Using 2 again,

$$\begin{aligned} Q_t^\ell (P_t^\ell - \alpha Q_t^\ell) &= \left\langle q_0 1^* - \ell \sqcup (1), \widehat{\mathbb{X}}_{0,t}^{<\infty} \right\rangle \left\langle e_{(2)} + 1^* - g^\ell - \alpha(q_0 1^* - \ell \sqcup (1)), \widehat{\mathbb{X}}_{0,t}^{<\infty} \right\rangle \\ &= \left\langle (q_0 1^* - \ell \sqcup (1)) \sqcup\!\sqcup (e_{(2)} + 1^* - g^\ell) - \alpha (q_0 1^* - \ell \sqcup (1))^{\sqcup\!\sqcup 2}, \widehat{\mathbb{X}}_{0,t}^{<\infty} \right\rangle . \end{aligned}$$

■

Acknowledgments. The authors would like to thank Sebastian Jaimungal, Alvaro Cartea, and Leandro Sanchez Betancourt for reading a preprint of this paper and for giving their insights.

Disclosure statement. Opinions and estimates constitute our judgement as of the date of this Material, are for informational purposes only, and are subject to change without notice. This Material is not the product of J.P. Morgans Research Department and therefore has not been prepared in accordance with legal requirements to promote the independence of research, including but not limited to the prohibition on the dealing ahead of the dissemination of investment research. This Material is not intended as research, a recommendation, advice, offer or solicitation for the purchase or sale of any financial product or service, or to be used in any way for evaluating the merits of participating in any transaction. It is not a research report and is not intended as such. Past performance is not indicative of future results. Please consult your own advisors regarding legal, tax, accounting, or any other aspects including suitability implications for your particular circumstances. J.P. Morgan disclaims any responsibility or liability whatsoever for the quality, accuracy, or completeness of the information herein and for any reliance on or use of this material in any way.

Important disclosures at www.jpmorgan.com/disclosures.

REFERENCES

[1] R. Almgren and N. Chriss, *Optimal execution of portfolio transactions*, J. Risk, 3 (2001), pp. 5–40.
[2] A. Ananova and R. Cont, *Pathwise integration with respect to paths of finite quadratic variation*, J. Math. Pures Appl. (9), 107 (2017), pp. 737–757.
[3] E. Bacry, A. Iuga, M. Lasnier, and C.-A. Lehalle, *Market impacts and the life cycle of investors orders*, Market Microstruct. Liquidity, 1 (2015), p. 1550009.
[4] M. Beiglböck, A. M. Cox, M. Huesmann, N. Perkowski, and D. J. Prömel, *Pathwise superreplication via Vovk's outer measure*, Finance Stoch., 21 (2017), pp. 1141–1166.
[5] D. Bertsimas and A. W. Lo, *Optimal control of execution costs*, J. Financ. Markets, 1 (1998), pp. 1–50.
[6] H. Boedihardjo, X. Geng, T. Lyons, and D. Yang, *The signature of a rough path: Uniqueness*, Adv. Math., 293 (2016), pp. 720–737.
[7] V. I. Bogachev, *Measure Theory*, Springer, Berlin, 2007.
[8] Á. Cartea, R. F. Donnelly, and S. Jaimungal, *Algorithmic trading with model uncertainty*, SIAM J. Financial Math., 8 (2017), pp. 635–671.
[9] Á. Cartea and S. Jaimungal, *Optimal execution with limit and market orders*, Quant. Finance, 15 (2015), pp. 1279–1291.

[10] Á. Cartea and S. Jaimungal, *A closed-form execution strategy to target volume weighted average price*, SIAM J. Financial Math., 7 (2016), pp. 760–785.
[11] Á. Cartea and S. Jaimungal, *Incorporating order-flow into optimal execution*, Math. Financ. Econ., 10 (2016), pp. 339–364.
[12] Á. Cartea, S. Jaimungal, and J. Penalva, *Algorithmic and High-Frequency Trading*, Cambridge University Press, Cambridge, UK, 2015.
[13] I. Chevyrev and T. J. Lyons, *Characteristic functions of measures on geometric rough paths*, Ann. Probab., 44 (2016), pp. 4049–4082.
[14] I. Chevyrev and H. Oberhauser, *Signature Moments to Characterize Laws of Stochastic Processes*, e-print, https://arxiv.org/abs/1810.10971, 2018.
[15] R. Cont and D.-A. Fourni, *Functional Itô calculus and stochastic integral representation of martingales*, Ann. Probab., 41 (2013), pp. 109–133.
[16] L. Coutin and A. Lejay, *Semi-martingales and rough paths theory*, Electron. J. Probab., 10 (2005), pp. 761–785.
[17] L. Coutin and Z. Qian, *Stochastic analysis, rough path analysis and fractional Brownian motions*, Probab. Theory Related Fields, 122 (2002), pp. 108–140.
[18] G. Curato, J. Gatheral, and F. Lillo, *Optimal execution with non-linear transient market impact*, Quant. Finance, 17 (2017), pp. 41–54.
[19] N.-M. Dang, *Optimal execution with transient impact*, Market Microstruct. Liquidity, 3 (2017), 1750008.
[20] B. Dupire, *Functional Itô Calculus*, Bloombery Portfolio research paper, 2009.
[21] G. Flint, B. Hambly, and T. Lyons, *Discretely sampled signals and the rough Hoff process*, Stochastic Process. Appl., 126 (2016), pp. 2593–2614.
[22] P. Friz and N. Victoir, *A note on the notion of geometric rough paths*, Probab. Theory Related Fields, 136 (2006), pp. 395–416.
[23] P. K. Friz and N. B. Victoir, *Multidimensional Stochastic Processes as Rough Paths: Theory and Applications*, Cambridge University Press, Cambridge, UK, 2010.
[24] J.-F. Le Gall, *A path-valued Markov process and its connections with partial differential equations*, in First European Congress of Mathematics, Paris, July 6–10, 1992, vol. 2, A. Joseph, F. Mignot, F. Murat, B. Prum, and R. Rentschler, eds., Birkhaüser, Basel, 1994, pp. 185–212.
[25] J. Gatheral, A. Schied, and A. Slynko, *Transient linear price impact and Fredholm integral equations*, Math. Finance, 22 (2012), pp. 445–474.
[26] C.-A. Lehalle and E. Neuman, *Incorporating signals into optimal trading*, Finance Stoch., 23 (2019), pp. 275–311.
[27] T. J. Lyons, *Differential equations driven by rough signals*, Rev. Mat. Iberoam., 14 (1998), pp. 215–310.
[28] T. J. Lyons, M. Caruana, and T. Lévy, *Differential Equations Driven by Rough Paths*, Springer, Berlin, 2007.
[29] I. Mastromatteo, M. Benzaquen, Z. Eisler, and J.-P. Bouchaud, *Trading Lightly: Cross-Impact and Optimal Portfolio Execution*, e-print, https://arxiv.org/abs/1702.03838, 2017.
[30] P. Pasquariello and C. Vega, *Strategic cross-trading in the U.S. stock market*, Rev. Finance, 19 (2015), pp. 229–282.
[31] J. Reizenstein and B. Graham, Algorithm 1004: *The iisignature library: Efficient calculation of iterated-integral signatures and log signatures*, ACM Trans. Math. Softw., 46 (2020), 8.
[32] C. Riga, *A pathwise Approach to Continuous-Time Trading*, e-print, https://arxiv.org/abs/1602.04946, 2016.
[33] M. H. Stone, *The generalized Weierstrass approximation theorem*, Math. Mag., 21 (1948), pp. 237–254.
[34] B. Toth, Y. Lemperiere, C. Deremble, J. de Lataillade, J. Kockelkoren, and J.-P. Bouchaud, *Anomalous price impact and the critical nature of liquidity in financial markets*, Phys. Rev. X, 1 (2011), 021006.
[35] G. Tsoukalas, J. Wang, and K. Giesecke, *Dynamic portfolio execution*, Manag. Sci., 65 (2017), pp. 1949–2443.