

ARTICLE TEMPLATE

Numerical method for model-free pricing of exotic derivatives using rough path signatures

Terry Lyons^{a, b}, Sina Nejad^{a, b} and Imanol Perez Arribas^{a, b, c}

^aMathematical Institute, University of Oxford, UK; ^bThe Alan Turing Institute, London, UK; ^cJ.P. Morgan, London, UK

ARTICLE HISTORY

Compiled February 4, 2020

ABSTRACT

We estimate prices of exotic options in a discrete-time model-free setting when the trader has access to market prices of a rich enough class of exotic and vanilla options. This is achieved by estimating an unobservable quantity called “implied expected signature” from such market prices, which are used to price other exotic derivatives. The implied expected signature is an object that characterises the market dynamics.

KEYWORDS

model-free pricing; rough path theory; financial derivatives

1. Introduction

A framework that has often been considered in model-free finance (see [1, 10, 8, 22]) 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.

The objective of this paper is to extend the ideas proposed in [27] to discrete-time, model-free frameworks. In [27], the authors introduce a family of basic exotic derivatives (called *signature payoffs*) that are used to price (or hedge) other exotic options. In this paper, we extend this approach to obtain a numerical method to price exotic derivatives from market data in discrete time in settings with minimal

Opinions expressed in this paper are those of the authors, and do not necessarily reflect the view of JP Morgan.

The authors would like to thank Samuel Cohen and David Prömel for their helpful insights towards the improvement of this paper.

This work was supported by The Alan Turing Institute under the EPSRC grant EP/N510129/1.

CONTACT Imanol Perez Arribas. Email: imanol.perez@maths.ox.ac.uk.

assumptions about the market. We illustrate our methodology with various numerical experiments that demonstrate the feasibility of our approach.

It is well-known ([9]) 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 of the same maturity by writing such contracts as linear combinations of call and put options. Signature options extend the idea of approximating complex contracts in terms of simpler, basic contracts to the setting of exotic path-dependent payoffs.

Our approach is the following: 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 characterises the entire market dynamics. Finally, we use the inferred implied expected signature to price other exotic payoffs.

Exotic derivatives are typically illiquid and 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 Mercure² and Totem Markit³. 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.

This paper assumes 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 discounted 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, 10, 8]). This is called the pricing-hedging duality.

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 financial contracts. The expected signature of a stochastic process ([12, 28, 13, 29]) 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 ([13, 12]), just like under some assumptions the moments

²<http://www.derivexperts.com/services/mercure>

³<https://ihsmarkit.com/products/totem.html>

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. Expected signatures have also been used in other areas of mathematical finance, such as in optimal execution [24, 11].

The remainder of the paper is structured as follows. In Section 2 we introduce the notion of the signature of a path, a crucial object in this paper. Section 3 introduces the financial framework we will consider. In Section 4 we introduce signature payoffs and we show how they are used to numerically estimate prices of other exotic derivatives. Finally, in Section 5 we show the feasibility of our numerical method with numerical experiments. All proofs have been postponed to Appendix A.

2. Signature of a path

Consider the d -dimensional Euclidean space \mathbb{R}^d . Denote by $(\mathbb{R}^d)^{\otimes n} := \underbrace{\mathbb{R}^d \otimes \dots \otimes \mathbb{R}^d}_n$, where \otimes is the tensor product. We define the *extended tensor algebra* over \mathbb{R}^d as

$$T((\mathbb{R}^d)) := \{(a_0, a_1, \dots, a_n, \dots) : a_n \in (\mathbb{R}^d)^{\otimes n} \text{ for all } n \geq 0\}.$$

Similarly, we denote

$$T(\mathbb{R}^d) := \{(a_0, a_1, \dots, a_n, \dots) \in T((\mathbb{R}^d)) : \exists N \in \mathbb{N} \text{ s.t. } a_n = 0 \text{ for all } n \geq N\}, \quad (1)$$

$$T^N(\mathbb{R}^d) := \{(a_0, a_1, \dots, a_n, \dots) \in T((\mathbb{R}^d)) : a_n = 0 \text{ for all } n \geq N\} \quad (2)$$

which we call the *tensor algebra* and *truncated tensor algebra at $N \in \mathbb{N}$* , respectively. All three are algebras with the product \otimes and sum $+$.

Denote by $(\mathbb{R}^d)^*$ the dual space of \mathbb{R}^d . Define the projection $\mathbf{1}^* \in T((\mathbb{R}^d)^*)$ by $\langle \mathbf{1}^*, \mathbf{a} \rangle := a_0$ for all $\mathbf{a} = (a_0, a_1, \dots) \in T((\mathbb{R}^d))$. Let $\{e_1, \dots, e_d\}$ be a basis for \mathbb{R}^d , and $\{e_1^*, \dots, e_d^*\}$ the corresponding dual basis for $(\mathbb{R}^d)^*$. Then,

$$\{e_I^* := e_{i_1}^* \otimes \dots \otimes e_{i_n}^* \mid I = (i_1, \dots, i_n) \in \{1, \dots, d\}^n, n \in \mathbb{N}\} \cup \{\mathbf{1}^*\}$$

is a basis for $T((\mathbb{R}^d))^* \cong T((\mathbb{R}^d)^*)$. There is a certain product that can be defined on $T((\mathbb{R}^d)^*)$, called the *shuffle product*:

Definition 2.1. Let $I = (i_1, \dots, i_n) \in \{1, \dots, d\}^n$, $J = (j_1, \dots, j_m) \in \{1, \dots, d\}^m$. The shuffle product of e_I^* and e_J^* , denoted by \sqcup , is defined inductively by:

$$e_I^* \sqcup e_J^* := (e_{(i_1, \dots, i_{n-1})}^* \sqcup e_J^*) \otimes e_{i_n}^* + (e_I^* \sqcup e_{(j_1, \dots, j_{m-1})}^*) \otimes e_{j_m}^*,$$

and $e_I^* \sqcup \mathbf{1}^* = \mathbf{1}^* \sqcup e_I^* = e_I^*$. The operation is extended by bilinearity to $T((\mathbb{R}^d)^*)$.

The signature of a path with bounded variation is a $T((\mathbb{R}^d))$ -valued object, defined as follows:

Definition 2.2 (Signature of a path). Given a path $Z \in C([0, T]; \mathbb{R}^d)$ with bounded variation, we denote its signature on $[s, t] \subset [0, T]$ by

$$\text{Sig}(Z)_{s,t} := (1, Z_{s,t}^1, Z_{s,t}^2, \dots, Z_{s,t}^n, \dots) \in T((\mathbb{R}^d))$$

where

$$Z_{s,t}^n := \int_{s < u_1 < \dots < u_n < t} dZ_{u_1} \otimes \dots \otimes dZ_{u_n} \in (\mathbb{R}^d)^{\otimes n}. \quad (3)$$

Similarly, the truncated signature of order $N \in \mathbb{N}$ is denoted by

$$\text{Sig}^N(Z)_{s,t} := (1, Z_{s,t}^1, \dots, Z_{s,t}^N) \in T^N(\mathbb{R}^d).$$

When we refer to the *signature of Z* without specifying the interval $[s, t]$, we will implicitly refer to the signature on $[0, T]$, i.e. $\text{Sig}(Z)_{0,T}$.

A thorough study of the signature of a path is beyond the scope of this paper – we refer the reader to [26] for a more detailed view of signatures and rough paths. We state the following key result; it guarantees that linear functions on the signature form an algebra.

Proposition 2.3. *[[26, Theorem 2.15]] Let $Z : [0, T] \rightarrow \mathbb{R}^d$ be a continuous path with bounded variation. Let $\ell_1, \ell_2 \in T((\mathbb{R}^d)^*)$ be two linear functionals. Then,*

$$\langle \ell_1, \text{Sig}(Z)_{s,t} \rangle \langle \ell_2, \text{Sig}(Z)_{s,t} \rangle = \langle \ell_1 \sqcup \ell_2, \text{Sig}(Z)_{s,t} \rangle \quad \forall 0 \leq s < t \leq T.$$

3. Framework

3.1. The market

Let $T > 0$ and let $\mathbb{T} = \{t_i\}_{i=0}^n \subset [0, 1]$ with $0 = t_0 < \dots < t_n = 1$ 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 $\Omega := \{X : \mathbb{T} \rightarrow \mathbb{R}_+ : X_0 = 1\}$. Therefore, we assume that the initial price is normalised to 1.




Figure 1.: Lead-lag transformation of a price path. 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.

Traders make decisions based on past information and the outcome of these decisions is known as new information is revealed. 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 [19] and is defined below.

Definition 3.1 (Lead-lag transformation). Given a price path $X \in \Omega$, define its lead-lag transformation $\hat{X} : [0, 1] \rightarrow \mathbb{R}^2 \oplus \mathbb{R}$ by the continuous path given by

$$\begin{aligned}\hat{X}_{2k/2n} &:= ((t_k, X_{t_k}), X_{t_k}) \in \mathbb{R}^2 \oplus \mathbb{R}, \\ \hat{X}_{(2k+1)/2n} &:= ((t_k, X_{t_k}), X_{t_{k+1}}) \in \mathbb{R}^2 \oplus \mathbb{R},\end{aligned}$$

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

Figure 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 $\hat{\Omega}$ are piecewise linear, $(\hat{\Omega}, \|\cdot\|_{BV})$ is separable. We consider the measure space $(\hat{\Omega}, \mathcal{B}(\hat{\Omega}))$.

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

Lemma 3.2 (Uniqueness of signature, [21]). *The signature map $\text{Sig} : \hat{\Omega} \rightarrow T((\mathbb{R}^2 \oplus \mathbb{R}))$ is injective.*

This result is not true if we consider the signature map $\text{Sig} : \Omega \rightarrow T((\mathbb{R}))$. In other words, given a path $X \in \Omega$, its signature $\text{Sig}(X)$ does not uniquely determine X itself, but if its lead-lag transformation $\hat{X} \in \hat{\Omega}$ is considered, then $\text{Sig}(\hat{X})$ *does* determine uniquely \hat{X} (and hence X). This is the main reason why lead-lag transformed paths are considered in this paper.

3.2. Payoffs and no-arbitrage assumptions

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

Definition 3.3 (Payoff). A payoff is a Borel-measurable function $G : \hat{\Omega} \rightarrow \mathbb{R}$.

The following are the assumptions we impose on the market.

Assumption 3.4. 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 $\mathfrak{F} = \{F_i\}_i$ that can be bought at $t = 0$, where $F_i : \hat{\Omega} \rightarrow \mathbb{R}$.
- (2) There are no market frictions such as transaction costs, slippage and bid-ask spread.
- (3) $\mathcal{M} \neq \emptyset$.
- (4) \mathcal{M} is tight.
- (5) $\mathbb{E}^{\mathbb{Q}}[\text{Sig}^N(\hat{X})_{0,1}]$ is well-defined for all $\mathbb{Q} \in \mathcal{M}$ and $N \in \mathbb{N}$.

The assumption that $\mathbb{E}^{\mathbb{Q}}[\text{Sig}^N(\hat{X})_{0,1}]$ 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 3.5. We denote by \mathcal{P} the pricing map, given by $\mathcal{P}(G) := \sup_{\mathbb{Q} \in \mathcal{M}} \mathbb{E}^{\mathbb{Q}}[G(\hat{X})]$ for each payoff G .

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

3.3. 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 [14, 2, 25, 4] and in [18, 7, 31] in the context of finance.

Definition 3.6 (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 3.1. This is a metric space with a certain metric ([14, Section 2.2]). 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 1. The space Λ is a Polish space (see [25]).

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 3.7. 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 paper we allow semi-static and dynamic hedging. Recall that \mathfrak{F} is the family of payoffs that can be bought at $t = 0$ (see Assumption 3.4). The performance of the hedging strategy will be determined by three factors: the price that is paid for each of the payoffs in \mathfrak{F} , the gains or losses that stem from holding the payoffs in \mathfrak{F} , and the gains that come from dynamically trading the underlying asset.

Definition 3.8. A hedging strategy is a pair $h = ((a_F)_{F \in \mathfrak{F}}, \theta)$, where $a_F \in \mathbb{R}$ for all $F \in \mathfrak{F}$ and $\theta \in \mathcal{T}$. The Profit and Loss (P&L) of h is defined by

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

We denote by \mathcal{H} the space of hedging strategies. A payoff $G : \widehat{\Omega} \rightarrow \mathbb{R}$ 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 $\varepsilon > 0$, $\mathcal{A}_\varepsilon := \{G : \widehat{\Omega} \rightarrow \mathbb{R} \text{ payoff} \mid \exists H \in \mathcal{A} \text{ such that } \|G - H\|_{L^\infty(\widehat{\Omega})} < \varepsilon\}$ the space of almost attainable payoffs.

4. Signature payoffs

We will now define a special class of continuous payoffs – the family of signature payoffs. These were first introduced in [27, 30], and as Arrow-Debreu securities [3, 15], they are *primitive securities* in the sense that path-dependent exotic derivatives can be approximated by these signature payoffs.

Definition 4.1 (Signature payoff). Let $\ell \in T((\mathbb{R}^2 \oplus \mathbb{R})^*)$. The signature payoff \mathcal{S}^ℓ is defined by

$$\begin{aligned}\mathcal{S}^\ell : \widehat{\Omega} &\rightarrow \mathbb{R} \\ \widehat{X} &\mapsto \langle \ell, \text{Sig}(\widehat{X})_{0,1} \rangle.\end{aligned}$$

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.

The following Proposition, see [27], shows that signature payoffs are primitive securities. Indeed, the proposition establishes that continuous payoffs can be locally approximated by signature payoffs:

Proposition 4.2. *Let $G : \widehat{\Omega} \rightarrow \mathbb{R}$ be a continuous payoff, and let $\mathcal{K} \subset \widehat{\Omega}$ be compact. Given any $\varepsilon > 0$, there exists a signature payoff \mathcal{S}^ℓ 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. *Let $G : \widehat{\Omega} \rightarrow \mathbb{R}$ be a continuous payoff. Let $\varepsilon > 0$. Then, there exists a compact $\mathcal{K}_\varepsilon \subset \widehat{\Omega}$ and a signature payoff \mathcal{S}^ℓ 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 4.4. *Set $\varepsilon > 0$. Let $G : \widehat{\Omega} \rightarrow \mathbb{R}$ 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}^ℓ such that, for $\mathcal{L} := \mathbf{1}_{\mathcal{K}_\varepsilon} \mathcal{S}^\ell : \widehat{\Omega} \rightarrow \mathbb{R}$,*

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

$$(2) \quad \mathcal{P}(|G - \mathcal{L}|) = \sup_{\mathbb{Q} \in \mathcal{M}} \mathbb{E}^{\mathbb{Q}}[|G(\hat{X}) - \mathcal{L}(\hat{X})|] < \varepsilon.$$

4.1. The implied expected signature

Let $\varepsilon > 0$. By Proposition 4.4, we may pick a compact set $\mathcal{K}_\varepsilon \subset \widehat{\Omega}$ and a family of signature payoffs $\mathcal{L} = \{\mathcal{S}^{\ell_F}\}_{F \in \mathfrak{F}}$ with $\ell_F \in T((\mathbb{R}^2 \oplus \mathbb{R})^*)$, such that:

- (1) $\mathbb{Q}(\hat{X} \in \mathcal{K}_\varepsilon) > 1 - \varepsilon$ for all $\mathbb{Q} \in \mathcal{M}$,
- (2) $\mathcal{P}(|F - \mathbb{1}_{\mathcal{K}_\varepsilon} \mathcal{S}^{\ell_F}|) < \varepsilon$ for all $F \in \mathfrak{F}$.

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

$$\langle \ell_F, \mathbb{E}^{\mathbb{Q}}[\text{Sig}(\hat{X})_{0,1} ; \mathcal{K}_\varepsilon] \rangle \approx \mathcal{P}(F) \quad \forall F \in \mathfrak{F}, \quad (4)$$

where $\mathbb{E}^{\mathbb{Q}}[\text{Sig}(\hat{X})_{0,1} ; \mathcal{K}_\varepsilon]$ denotes $\mathbb{E}^{\mathbb{Q}}[\text{Sig}(\hat{X})_{0,1} \mathbb{1}_{\mathcal{K}_\varepsilon}]$. Choose $N \in \mathbb{N}$ sufficiently large so that $\ell_F \in T^N(\mathbb{R}^2 \oplus \mathbb{R})$ for all $F \in \mathfrak{F}$, which is possible because \mathfrak{F} is finite and by definition of tensor algebra (1). Then, (4) becomes:

$$\langle \ell_F, \mathbb{E}^{\mathbb{Q}}[\text{Sig}^N(\hat{X})_{0,1} ; \mathcal{K}_\varepsilon] \rangle \approx \mathcal{P}(F) \quad \forall F \in \mathfrak{F}. \quad (5)$$

If the risk-neutral measures are unknown and we only observe market prices for the exotic payoffs \mathfrak{F} , we may look for the expected signature that best matches the observed prices – i.e. the best expected signature fit for the approximation (5). This can be achieved using linear regression. The complexity of the estimation will, in general, depend on the dimension of $T^N(\mathbb{R}^2 \oplus \mathbb{R})$ – which is $(3^{N+1} - 1)/2$ – and the number of exotics options $|\mathfrak{F}|$. Notice, in particular, that the complexity grows exponentially with N . In the experiments in Section 5, however, it sufficed to set $N = 5$.

The extent to which we can accurately estimate the expected signature depends on the invertibility of the linear map

$$\begin{aligned} T^N(\mathbb{R}^2 \oplus \mathbb{R}) &\rightarrow \mathbb{R}^{|\mathfrak{F}|} \\ \mathbf{a} &\mapsto (\langle \ell_F, \mathbf{a} \rangle)_{F \in \mathfrak{F}}. \end{aligned} \quad (6)$$

If the family of exotic payoffs \mathfrak{F} is rich enough (i.e. the linear map (6) 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 concept of *implied volatility*.

If we are given a (potentially illiquid) payoff $G : \widehat{\Omega} \rightarrow \mathbb{R}$ that is close to the space of attainable payoffs \mathcal{A} (Definition 3.8) and that satisfies the assumptions of Proposition 4.4, we may estimate its price by $\mathcal{P}(G) \approx \langle \ell_G, \mathbf{E} \rangle$ where ℓ_G is given by Proposition 4.4

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

This approach provides a numerical 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 \mathfrak{F} with observable prices.

Example 4.5. In what follows, we discuss a couple of examples to provide intuition of the meaning of some of the terms of the implied expected signature.

Let $\hat{X} = (X^b, X^f) \in \hat{\Omega}$ and write $\hat{X}_t = ((\hat{X}_t^1, \hat{X}_t^2), \hat{X}_t^3) \in \mathbb{R}^2 \oplus \mathbb{R}$. By the definition of the iterated integrals (3) that form the signature, the first term of the signature is given by:

$$\langle e_2^*, \text{Sig}(\hat{X})_{0,1} \rangle \stackrel{\text{def}}{=} \int_0^1 d\hat{X}_t^1 = X_1 - X_0.$$

In other words, the first term of the signature is given by the increment of the price path. Hence, the first term of the implied expected signature $\langle e_2^*, \mathbf{E} \rangle$ is an estimate of the forward price of the underlying asset.

For the second term of the signature, the lead-lag transformation \hat{X} has the property that

$$\left\langle e_3^* \otimes e_2^* - e_2^* \otimes e_3^*, \text{Sig}(\hat{X})_{0,1} \right\rangle \stackrel{\text{def}}{=} \int_0^1 \int_0^t d\hat{X}_s^3 d\hat{X}_t^2 - \int_0^1 \int_0^t d\hat{X}_s^2 d\hat{X}_t^3 = \langle X \rangle,$$

where $\langle X \rangle = \sum_{i=1}^n (X_{t_{i-1}} - X_{t_i})^2$ is the quadratic variation – or volatility – of the price path. Therefore, the second term of the implied expected signature $\langle e_3^* \otimes e_2^* - e_2^* \otimes e_3^*, \mathbf{E} \rangle$ is an estimate of the expected volatility of the market.

In a nutshell, the first two order terms of the implied expected signature contain information about the forward price and expected volatility of the market. Higher order terms capture other aspects of the market dynamics. In fact, the full implied expected signature (when $N \rightarrow \infty$) characterises the market dynamics, see [13].

5. Numerical experiment

5.1. Approximating payoffs by signature payoffs

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

We begin by choosing a dataset of paths $\{\hat{X}^{(i)}\}_{i=1}^n \subset \hat{\Omega}$. In the experiments in this section, we generated 10,000 paths from a Black–Scholes model with a constant volatility ranging from 5% to 40% – 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 $\{\text{Sig}^N(\hat{X}^{(i)})_{0,1}\}_{i=1}^n$ 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`⁴ or `iisignature`⁵.

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

$$\langle \ell, \text{Sig}^N(\hat{X}^{(i)})_{0,1} \rangle \approx G(\hat{X}^{(i)}) \quad \forall i = 1, \dots, n.$$

5.2. Pricing from market data

In this section we implement the proposed approach to estimate the implied expected signature from a family of exotic options \mathfrak{F} , which we employ to price other exotic options. The family \mathfrak{F} of exotic options available to buy at $t = 0$ (see Assumption 3.4) that was considered consists of 100 exotic and vanilla payoffs with maturity 1 year. More specifically, we consider European options, barrier up-and-out options, barrier up-and-in options and variance options (25 payoffs of each type). The size of the dataset, as well as 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 [23], GARCH model [5, 6] and the rough volatility model [20]. 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 5.1 to approximate each payoff in \mathfrak{F} by a signature payoff, and then we followed the numerical method proposed in Section 4.1 to estimate the implied expected signature. Recall that 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 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 4.1. 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 3, where the R^2 of the predicted market prices of exotic payoffs is shown as a function

⁴<https://pypi.org/project/esig/>


⁵<https://github.com/bottler/iisignature>



(a) Hull-White model.



(b) GARCH model.



(c) Rough volatility model.

Figure 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.




Figure 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} $ \mathfrak{F} = 100$	Vanilla and exotic options in \mathfrak{F}		
		$ \mathfrak{F} = 33$	$ \mathfrak{F} = 66$	$ \mathfrak{F} = 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}	1.48×10^{-8}
Mean absolute error	2.94×10^{-2}	7.89×10^{-2}	6.04×10^{-4}	3.39×10^{-5}

Table 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 options, barrier up-and-in options and variance option). The table suggests that a large family of payoffs \mathfrak{F} that includes exotic options leads to the best performance.

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} (i.e. the set of traded derivatives) 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 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 3, where the accuracy of predicted market prices of exotic payoffs is shown in the case where \mathfrak{F} only contains vanilla options.

Table 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.

6. Conclusion

In this paper we proposed a numerical method for estimating prices of exotic derivatives from available market prices of other exotic derivatives. We do so in a model-free setting, where only minimal assumptions about the market are made.

Our methodology extends the ideas from the contemporaneous work in [27]. We show that from market prices of exotic derivatives we can estimate a certain non-observable quantity called the *implied expected signature*. This object consists of a collection of statistics that characterise the market dynamics, and it can be used to numerically estimate market prices of other exotic derivatives.

We believe there are several directions that this paper can be extended in. For instance, core ideas from the paper could be applied to risk management of large portfolios of exotic derivatives: portfolio optimisation, value at risk (VaR) calculation, etc. Further research in this direction may lead to new tools in these fields.

Appendix A. Proofs

Proof of Proposition 4.2. First, we show that the class of all signature payoffs forms an algebra. Let $\ell_1, \ell_2 \in T((\mathbb{R}^2 \oplus \mathbb{R})^*)$, and consider the corresponding signature payoffs $\mathcal{S}^{\ell_1}, \mathcal{S}^{\ell_2} : \widehat{\Omega} \rightarrow \mathbb{R}$. 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 (see Proposition 2.3), we have:

$$\mathcal{S}^{\ell_1} \cdot \mathcal{S}^{\ell_2} = \mathcal{S}^{\ell_1 \sqcup \ell_2}.$$

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 3.2, the signature map $\widehat{\Omega} \rightarrow T((\mathbb{R}^2 \oplus \mathbb{R}))$ is injective. Therefore, given $\widehat{X}_1, \widehat{X}_2 \in \widehat{\Omega}$ distinct, we have $\text{Sig}(\widehat{X}_1) \neq \text{Sig}(\widehat{X}_2)$. It follows immediately that there exists a signature payoff \mathcal{S}^ℓ , 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 $G : \widehat{\Omega} \rightarrow \mathbb{R}$ and a compact set $\mathcal{K} \subset \widehat{\Omega}$, there exists a signature payoff \mathcal{S}^ℓ with $\ell \in T((\mathbb{R}^2 \oplus \mathbb{R})^*)$ such that $|G(\widehat{X}) - \mathcal{S}^\ell(\widehat{X})| < \varepsilon$ for all $\widehat{X} \in \mathcal{K}$. \square

Proof of Corollary 4.3. 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 4.2 with the payoff G and compact \mathcal{K}_ε , and conclude that $\|G - \mathcal{S}^\ell\|_{L^\infty(\mathcal{K}_\varepsilon)} < \varepsilon$. \square

Proof of Proposition 4.4. 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, 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

$$\begin{aligned} \mathcal{P}(|G - \mathcal{L}|) &= \sup_{\mathbb{Q} \in \mathcal{M}} \mathbb{E}^\mathbb{Q}[|G(\widehat{X}) - \mathcal{L}(\widehat{X})|] = \sup_{\mathbb{Q} \in \mathcal{M}} (\mathbb{E}^\mathbb{Q}[|G(\widehat{X})| ; \mathcal{K}_\varepsilon^c] + \mathbb{E}^\mathbb{Q}[|G(\widehat{X}) - \mathcal{S}^\ell(\widehat{X})| ; \mathcal{K}_\varepsilon]) \\ &< \varepsilon. \end{aligned}$$

□

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] Acciaio, B., Beiglbck, M., Penkner, F. and Schachermayer, W., 2016. A model-free version of the fundamental theorem of asset pricing and the superreplication theorem. *Mathematical Finance*, 26(2), pp.233-251.
- [2] Ananova, A. and Cont, R., 2017. Pathwise integration with respect to paths of finite quadratic variation. *Journal de Mathématiques Pures et Appliquées*, 107(6), pp.737-757.
- [3] Arrow, K.J., 1973. The role of securities in the optimal allocation of risk-bearing. In *Readings in Welfare Economics* (pp. 258-263). Palgrave, London.
- [4] Bally, V., Caramellino, L., Cont, R., i Civit, F.U. and Vives, J., 2016. *Stochastic integration by parts and functional It calculus*. Birkhuser.
- [5] BarndorffNielsen, O.E. and Shephard, N., 2001. NonGaussian OrnsteinUhlenbeck-based models and some of their uses in financial economics. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 63(2), pp.167-241.
- [6] Barndorff-Nielsen, O.E. and Shephard, N., 2001. Modelling by Lvy processes for financial econometrics. In *Lvy processes* (pp. 283-318). Birkhuser, Boston, MA.
- [7] Beiglbck, M., Cox, A.M., Huesmann, M., Perkowski, N. and Prmel, D.J., 2017. Pathwise superreplication via Vovks outer measure. *Finance and Stochastics*, 21(4), pp.1141-1166.
- [8] Bouchard, B. and Nutz, M., 2015. Arbitrage and duality in nondominated discrete-time models. *The Annals of Applied Probability*, 25(2), pp.823-859.
- [9] Breeden, D.T. and Litzenberger, R.H., 1978. Prices of state-contingent claims implicit in option prices. *Journal of business*, pp.621-651.
- [10] Burzoni, M., Frittelli, M., Hou, Z., Maggis, M. and Oblj, J., 2016. Pointwise arbitrage pricing theory in discrete time. *arXiv preprint arXiv:1612.07618*.
- [11] Cartea, Á., Pérez Arribas, I. and Sánchez-Betancourt, L., 2020. Optimal Execution of Foreign Securities: A Double-Execution Problem with Machine Learning. *Working paper*.
- [12] Chevyrev, I. and Oberhauser, H., 2018. Signature moments to characterize laws of stochastic processes. *arXiv preprint arXiv:1810.10971*.
- [13] Chevyrev, I. and Lyons, T., 2016. Characteristic functions of measures on geometric rough paths. *The Annals of Probability*, 44(6), pp.4049-4082.
- [14] Cont, R. and Fourni, D.A., 2013. Functional It calculus and stochastic integral representation of martingales. *The Annals of Probability*, 41(1), pp.109-133.
- [15] Debreu, G., 1987. *Theory of value: An axiomatic analysis of economic equilibrium* (No. 17). Yale University Press.
- [16] Derman, E. and Kani, I., 1994. Riding on a smile. *Risk*, 7(2), pp.32-39.
- [17] Dupire, B., 1994. Pricing with a smile. *Risk*, 7(1), pp.18-20.
- [18] Dupire, B., 2009. Functional It calculus.
- [19] Flint, G., Hambly, B. and Lyons, T., 2016. Discretely sampled signals and the

- rough Hoff process. *Stochastic Processes and their Applications*, 126(9), pp.2593-2614.
- [20] Gatheral, J., Jaisson, T. and Rosenbaum, M., 2014. Volatility is rough. *arXiv preprint arXiv:1410.3394*.
 - [21] Hambly, B. and Lyons, T., 2010. Uniqueness for the signature of a path of bounded variation and the reduced path group. *Annals of Mathematics*, pp.109-167.
 - [22] Hou, Z. and Obłój, J., 2018. Robust pricing-hedging dualities in continuous time. *Finance and Stochastics*, 22(3), pp.511-567.
 - [23] Hull, J. and White, A., 1987. The pricing of options on assets with stochastic volatilities. *The journal of finance*, 42(2), pp.281-300.
 - [24] Kalsi, J., Lyons, T. and Pérez Arribas, I., 2019. Optimal execution with rough path signatures. *arXiv preprint arXiv:1905.00728*.
 - [25] Le Gall, J.F., 1994. A path-valued Markov process and its connections with partial differential equations. In *First European Congress of Mathematics Paris, July 610*, 1992 (pp. 185-212). Birkhäuser Basel.
 - [26] Lyons, T.J., Caruana, M. and Lévy, T., 2007. Differential equations driven by rough paths. Springer Berlin Heidelberg.
 - [27] Lyons, T., Nejad, S. and Pérez Arribas, I., 2019. Nonparametric pricing and hedging of exotic derivatives. *arXiv preprint arXiv:1905.00711*.
 - [28] Lyons, T. and Ni, H., 2015. Expected signature of Brownian motion up to the first exit time from a bounded domain. *The Annals of Probability*, 43(5), pp.2729-2762.
 - [29] Ni, H., 2012. *The expected signature of a stochastic process* (Doctoral dissertation, Oxford University, UK).
 - [30] Pérez Arribas, I., 2018. Derivatives pricing using signature payoffs. *arXiv preprint arXiv:1809.09466*.
 - [31] Riga, C., 2016. A pathwise approach to continuous-time trading. *arXiv preprint arXiv:1602.04946*.