## Appendix

This Appendix is organised as follows: A) using the formalism of cross-covariance operators we define an Hilbert-Schmidt conditional independence criterion for stochastic processes, and provide further details on the construction of the estimator for the $2^{\text {nd }}$ order MMD; B) we outline algorithms and their complexities for computing higher order MMDs; C) we provide further experimental details; D) we prove the theorems from the main paper.

## A Cross-covariance operators

Covariance and cross-covariance operators on RKHSs are important concepts for modern applications of conditional KMEs [50, 10]. In this section we will use this formalism (adapted to the case of path-valued random variables) to firstly derive a criterion for conditional independence of stochastic processes and secondly provide more details on the derivation of our estimator of the $2^{\text {nd }}$ order MMD.

Let $X, Y \in \mathcal{P}(\mathcal{X}(V))$ be two stochastic processes and their joint process $(X, Y) \in \mathcal{P}(\mathcal{X}(V \oplus V))$. Define the cross-covariance operator $\mathcal{C}_{Y, X}$ as the following point in the tensor product of RKHSs $\mathcal{H}_{\mathcal{S}}(V) \otimes \mathcal{H}_{\mathcal{S}}(V)$

$$
\begin{equation*}
\mathcal{C}_{Y, X}=\mathbb{E}_{(X, Y)}\left[k_{\mathcal{S}}(\cdot, Y) \otimes k_{\mathcal{S}}(X, \cdot)\right] \tag{28}
\end{equation*}
$$

or equivalently as the Hilbert-Schmidt operator $\mathcal{C}_{Y, X}: \mathcal{H}_{\mathcal{S}}(V) \rightarrow \mathcal{H}_{\mathcal{S}}(V)$ defined for any function $f \in \mathcal{H}_{\mathcal{S}}(V)$ as follows

$$
\begin{equation*}
\mathcal{C}_{Y, X}(f)(\cdot)=\int_{(x, y) \in \mathcal{X}(V \oplus V)} k_{\mathcal{S}}(\cdot, y) f(x) \mathbb{P}_{(X, Y)}(d(x, y)) \tag{29}
\end{equation*}
$$

The equivalence between tensor product of RKHSs and Hilbert-Schmidt operators is given by the isomorphism $\Phi: \mathcal{H}_{\mathcal{S}}(V) \otimes \mathcal{H}_{\mathcal{S}}(V) \rightarrow \operatorname{HS}\left(\mathcal{H}_{\mathcal{S}}(V), \mathcal{H}_{\mathcal{S}}(V)\right)$ defined as follows

$$
\begin{equation*}
\Phi_{\mathcal{S}}\left(\sum_{\mathbf{k}, \mathbf{k}^{\prime}} \alpha_{\mathbf{k}, \mathbf{k}^{\prime}} \mathcal{S}^{(\mathbf{k})} \otimes \mathcal{S}^{\left(\mathbf{k}^{\prime}\right)}\right)=\sum_{\mathbf{k}, \mathbf{k}^{\prime}} \alpha_{\mathbf{k}, \mathbf{k}^{\prime}}\left\langle\cdot, \mathcal{S}^{(\mathbf{k})}\right\rangle_{\mathcal{H}_{\mathcal{S}}(V)} \mathcal{S}^{\left(\mathbf{k}^{\prime}\right)} \tag{30}
\end{equation*}
$$

where $\mathcal{S}^{(\mathbf{k})}$ denotes the $\mathbf{k}^{\text {th }}$ element of an orthogonal basis of $\mathcal{H}_{\mathcal{S}}(V)$ and $\operatorname{HS}\left(\mathcal{H}_{\mathcal{S}}(V), \mathcal{H}_{\mathcal{S}}(V)\right)$ is the space of Hilbert-Schmidt operators from $\mathcal{H}_{\mathcal{S}}(V)$ to itself. An example of such basis is given by the signature basis defined for any path $x \in \mathcal{X}(V)$ and any coordinate $\mathbf{k}=\left(k_{1}, \ldots, k_{j}\right)$ as

$$
\begin{equation*}
\mathcal{S}^{\mathbf{k}}(x)=\int_{0<s_{1}<\ldots<s_{j}<T} \ldots \int_{s_{1}} d x_{s_{j}}^{\left(k_{1}\right)} \ldots d x_{s_{j}}^{\left(k_{j}\right)} \tag{31}
\end{equation*}
$$

The centered version $\tilde{\mathcal{C}}_{Y, X} \in \mathcal{H}_{\mathcal{S}}(V) \otimes \mathcal{H}_{\mathcal{S}}(V)$ of the operator $\mathcal{C}_{Y, X}$ is defined as

$$
\begin{equation*}
\tilde{\mathcal{C}}_{Y, X}=\mathcal{C}_{Y, X}-\mu_{X}^{1} \otimes \mu_{Y}^{1} \tag{32}
\end{equation*}
$$

Similarly let $\mathcal{C}_{X, X} \in \mathcal{H}_{\mathcal{S}}(V) \otimes \mathcal{H}_{\mathcal{S}}(V)$ be the following covariance operator

$$
\begin{equation*}
\mathcal{C}_{X, X}=\mathbb{E}_{X}\left[k_{\mathcal{S}}(X, \cdot) \otimes k_{\mathcal{S}}(X, \cdot)\right] \tag{33}
\end{equation*}
$$

or equivalently $\mathcal{C}_{X, X} \in \operatorname{HS}\left(\mathcal{H}_{\mathcal{S}}(V), \mathcal{H}_{\mathcal{S}}(V)\right)$

$$
\begin{equation*}
\mathcal{C}_{X, X}(f)(\cdot)=\int_{x \in \mathcal{X}(V)} k_{\mathcal{S}}(\cdot, x) f(x) \mathbb{P}_{X}(d x) \tag{34}
\end{equation*}
$$

Under the assumption that for any $f \in \mathcal{H}_{\mathcal{S}}(V)$ the function $x \mapsto \mu_{f(Y) \mid X=x}^{1}$ from $\mathcal{X}(V)$ to $\mathbb{R}$ is in $\mathcal{H}_{\mathcal{S}}(V)$, the authors in $[10,50]$ showed that

$$
\begin{equation*}
\mu_{Y \mid X}^{1}=\mathcal{C}_{Y, X} \mathcal{C}_{X, X}^{-1} \tag{35}
\end{equation*}
$$

However, this assumption might not hold in general [10,50]. This technical issue can be circumvented by resorting to a regularized version of eq. (35): this yields to

$$
\begin{equation*}
\mu_{Y \mid X}^{1} \approx \mathcal{C}_{Y, X}\left(\mathcal{C}_{X, X}+\lambda I_{\mathcal{H}_{\mathcal{S}}(V)}\right)^{-1}, \quad \lambda>0 \tag{36}
\end{equation*}
$$

where $I_{\mathcal{H}_{\mathcal{S}}(V)}$ is the identity map from $\mathcal{H}_{\mathcal{S}}(V)$ to itself. Under some mild conditions the empirical estimator of eq. (36) is equal to the empirical estimator of eq. (35) [9, Thm. 8]. In particular, one has

$$
\begin{equation*}
\hat{\mu}_{Y \mid X=x}^{1}=\mathbf{k}^{x \top}\left(\mathbf{K}^{x, x}+m \lambda I_{m}\right)^{-1} \mathbf{k}^{y}(\cdot) \tag{37}
\end{equation*}
$$

based on sample paths $\left\{\left(x^{i}, y^{i}\right)\right\}_{i=1}^{m}$ from the joint $(X, Y)$ where $\mathbf{k}^{x}, \mathbf{K}^{x, x}$ and $\mathbf{k}^{y}(\cdot)$ are such that

$$
\left[\mathbf{k}^{x}\right]_{i}=k_{\mathcal{S}}\left(x^{i}, x\right) \quad\left[\mathbf{K}^{x, x}\right]_{i, j}=k_{\mathcal{S}}\left(x^{i}, x^{j}\right) \quad \mathbf{k}^{y}(\cdot)=\left[k_{\mathcal{S}}\left(y_{1}, \cdot\right), \ldots, k_{\mathcal{S}}\left(y_{m}, \cdot\right)\right]^{\top}
$$

Cross-covariance operators have been used to define kernel-based measures of conditional dependence, as we shall discuss in the next section.

## A. 1 Hilbert-Schmidt Conditional Independence Criterion for stochastic processes

Multiple measures of conditional dependence have been proposed in the literature [14, 13, 12]. In this section, we follow [12] to define a nonparametric conditional dependence measure for stochastic processes, based on the conditional cross-covariance operator $\tilde{\mathcal{C}}_{Y, X \mid Z}: \mathcal{H}_{\mathcal{S}}(V) \rightarrow \mathcal{H}_{\mathcal{S}}(V)$,

$$
\begin{equation*}
\tilde{\mathcal{C}}_{Y, X \mid Z}=\tilde{\mathcal{C}}_{Y, X}-\tilde{\mathcal{C}}_{Y, Z} \tilde{\mathcal{C}}_{Z, Z}^{-1} \tilde{\mathcal{C}}_{Z, X} \tag{38}
\end{equation*}
$$

The squared Hilbert-Schmidt norm $H_{Y X \mid Z}:=\left\|\tilde{\mathcal{C}}_{(Y, Z), X \mid Z}\right\|_{\text {HS }}^{2}$ can be used as measure of conditional dependence of stochastic processes. Since the signature kernel $k_{\mathcal{S}}$ is characteristic, it follows that $X \Perp Y \mid Z \Longleftrightarrow H_{Y X \mid Z}=0$ [12].

Given $m$ sample paths $\left\{\left(x_{i}, y_{i}, z_{i}\right)\right\}_{i=1}^{m}$ from the joint distribution of $(X, Y, Z)$, let $K^{x}, K^{y}$ and $K^{z}$ be the Gram matrices with entries,

$$
\left[K^{x}\right]_{i, j}=k_{\mathcal{S}}\left(x_{i}, x_{j}\right) \quad\left[K^{y}\right]_{i, j}=k_{\mathcal{S}}\left(y_{i}, y_{j}\right) \quad\left[K^{z}\right]_{i, j}=k_{\mathcal{S}}\left(z_{i}, z_{j}\right)
$$

An empirical estimator of the kernel conditional dependence measure $H_{Y X \mid Z}$ is then given by,
$\widehat{H}_{Y X \mid Z}=\frac{1}{m^{2}}\left\{\operatorname{tr}\left(\tilde{K}^{x} \tilde{K}^{y}\right)-2 \operatorname{tr}\left(\tilde{K}^{x} \tilde{K}^{z}\left(\tilde{K}_{\epsilon}^{z}\right)^{-2} \tilde{K}^{z} \tilde{K}^{y}\right)+\operatorname{tr}\left(\tilde{K}^{x} \tilde{K}^{z}\left(\tilde{K}_{\epsilon}^{z}\right)^{-2} \tilde{K}^{z} \tilde{K}^{y} \tilde{K}^{z}\left(\tilde{K}_{\epsilon}^{z}\right)^{-2} \tilde{K}^{z}\right)\right\}$
where $\tilde{K}_{\epsilon}^{z}=\tilde{K}^{z}+\epsilon I_{m}$ and $\tilde{K}^{x}, \tilde{K}^{y}, \tilde{K}^{z}$ are the centered versions of the matrices $K^{x}, K^{y}$ and $K^{z}$,

$$
\tilde{K}^{x}=H K^{x} H \quad \tilde{K}^{y}=H K^{y} H \quad \tilde{K}^{z}=H K^{z} H
$$

with $H=I_{m}-m^{-1} \mathbf{1}_{m}$ and $\mathbf{1}_{m}$ the $m \times m$ matrix with all entries set to 1 . This estimator can be used as a test statistic for testing whether $X$ and $Y$ are independent given $Z$. However, it is not known how to analytically compute the null distribution of the test statistic, and permutation tests are typically used. In Sec. 4.3 we use this measure of conditional dependence as part of the kPC algorithm to infer causal relationships between multidimensional stochastic processes. We provide more details in Appendix C.

## A. 2 Construction of the estimator for the second order MMD $\mathcal{D}_{\mathcal{S}}^{2}$

As discussed in the main paper, the estimation of the $2^{\text {nd }}$ order MMD, required the ability to compute inner products of the form $\left\langle\tilde{x}_{s}, \tilde{y}_{t}\right\rangle$ in $\mathcal{H}_{\mathcal{S}}(V)$. Here, we provide more details on the approximation that we have used in eq. (20), also restated below,

$$
\left\langle\widetilde{x}_{s}, \widetilde{y}_{t}\right\rangle_{\mathcal{H}_{s}(V)} \approx \mathbf{k}_{s}^{x \top}\left(\mathbf{K}_{s, s}^{x, x}+m \lambda I_{m}\right)^{-1} \mathbf{K}_{T, T}^{x, y}\left(\mathbf{K}_{t, t}^{y, y}+n \lambda I_{n}\right)^{-1} \mathbf{k}_{t}^{y}
$$

where $\widetilde{x}$ and $\widetilde{y}$ are sample paths from the processes $\mu_{X \mid \mathcal{F}_{X}}^{1}$ and $\mu_{Y \mid \mathcal{F}_{Y}}^{1}$ respectively. In particular,

$$
\widetilde{x}_{s}=\mu_{X \mid x_{[0, s]}}^{1} \quad \text { and } \quad \widetilde{y}_{t}=\mu_{Y \mid y_{[0, t]}}^{1}
$$

As discussed at the beginning of this section, their empirical estimators are constructed from $m$ samples $\left\{x^{i}\right\}_{i=1}^{m}$ from $X$ and $n$ samples $\left\{y^{j}\right\}_{j=1}^{n}$ from $Y$ respectively

$$
\widetilde{x}_{s} \approx \mathbf{k}_{s}^{x^{\top}}\left(\mathbf{K}_{s, s}^{x, x}+m \lambda I_{m}\right)^{-1} \mathbf{k}^{x}(\cdot) \quad \text { and } \quad \widetilde{y}_{t} \approx \mathbf{k}_{t}^{y \top}\left(\mathbf{K}_{t, t}^{y, y}+n \lambda I_{n}\right)^{-1} \mathbf{k}^{y}(\cdot)
$$

where $\mathbf{k}_{s}^{x}, \mathbf{K}_{s, s}^{x, x}, \mathbf{k}^{x}(\cdot)$ and $\mathbf{k}_{t}^{y}, \mathbf{K}_{t, t}^{y, y}$ and $\mathbf{k}^{y}(\cdot)$ are defined by,

$$
\begin{array}{lll}
{\left[\mathbf{k}_{s}^{x}\right]_{i}=k_{\mathcal{S}}\left(x_{[0, s]}^{i}, x_{[0, s]}\right)} & {\left[\mathbf{K}_{s, s}^{x, x}\right]_{i, j}=k_{\mathcal{S}}\left(x_{[0, s]}^{i}, x_{[0, s]}^{j}\right)} & \mathbf{k}^{x}(\cdot)=\left[k_{\mathcal{S}}\left(x_{1}, \cdot\right), \ldots, k_{\mathcal{S}}\left(x_{m}, \cdot\right)\right]^{\top} \\
{\left[\mathbf{k}_{t}^{y}\right]_{i}=k_{\mathcal{S}}\left(y_{[0, t]}^{i}, y_{[0, t]}\right)} & {\left[\mathbf{K}_{t, t}^{y, y}\right]_{i, j}=k_{\mathcal{S}}\left(y_{[0, t]}^{i}, y_{[0, t]}^{j}\right)} & \mathbf{k}^{y}(\cdot)=\left[k_{\mathcal{S}}\left(y_{1}, \cdot\right), \ldots, k_{\mathcal{S}}\left(y_{n}, \cdot\right)\right]^{\top}
\end{array}
$$

Alternatively, we can write $\widetilde{x}_{s} \approx \sum_{i=1}^{m} \alpha_{i} k_{\mathcal{S}}\left(x^{i}, \cdot\right)$ with $\boldsymbol{\alpha}=\left(\mathbf{K}_{s, s}^{x, x}+m \lambda I_{m}\right)^{-1} \mathbf{k}_{s}^{x}$. Similarly we have $\widetilde{y}_{t} \approx \sum_{j=1}^{n} \beta_{j} k_{\mathcal{S}}\left(\cdot, y^{j}\right)$ with $\boldsymbol{\beta}=\left(\mathbf{K}_{t, t}^{y, y}+n \lambda I_{n}\right)^{-1} \mathbf{k}_{t}^{y}$. Therefore, the inner product between $\widetilde{x}_{s}$ and $\widetilde{y}_{t}$ can be approximated as follows,

$$
\begin{aligned}
\left\langle\widetilde{x}_{s}, \widetilde{y}_{t}\right\rangle_{\mathcal{H}(V)} & \approx \sum_{i=1}^{m} \sum_{j=1}^{n} \alpha_{i} \beta_{j} k_{\mathcal{S}}\left(x^{i}, y^{j}\right) \\
& =\boldsymbol{\alpha}^{\top} \mathbf{K}_{T, T}^{x, y} \boldsymbol{\beta} \\
& =\mathbf{k}_{s}^{x^{\top}}\left(\mathbf{K}_{s, s}^{x, x}+m \lambda I_{m}\right)^{-1} \mathbf{K}_{T, T}^{x, y}\left(\mathbf{K}_{t, t}^{y, y}+n \lambda I_{n}\right)^{-1} \mathbf{k}_{t}^{y}
\end{aligned}
$$

where $\mathbf{K}_{T, T}^{x, y} \in \mathbb{R}^{m \times n}$ with $\left[\mathbf{K}_{T, T}^{x, y}\right]_{i, j}=k_{\mathcal{S}}\left(x^{i}, y^{j}\right)$. Next we outline the algorithm to compute $\hat{\mathcal{D}}_{\mathcal{S}}^{2}$.

## B Algorithms

In this section, we provide algorithms to compute the empirical estimator $\hat{\mathcal{D}}_{\mathcal{S}}^{k}$ for the $k^{\text {th }}$ order MMD, which rely on the ability to evaluate the signature kernel $k_{\mathcal{S}}(x, y)$ where $x$ and $y$ are two paths taking their values in the Hilbert space $\mathcal{H}^{k-1}(V)$. Following [22, Sec. 3.1.] we use an explicit finite difference scheme to approximate the PDE solution $u_{x, y}$ on a grid $\mathcal{P}$ of size $P \times Q$,

$$
\mathcal{P}=\left\{0=s_{1}<s_{2}<\ldots<s_{P}=T\right\} \times\left\{0=t_{1}<t_{2}<\ldots<t_{Q}=T\right\}
$$

Writing $u_{x, y}\left(s_{i}, t_{j}\right)=u_{i, j}$ to make the notation more concise, we use an update rule of the form,

$$
u_{i+1, j+1}=f\left(u_{i, j+1}, u_{i+1, j}, u_{i, j}, M_{i, j}\right), \quad M_{i, j}=\left\langle x_{s_{i+1}}-x_{s_{i}}, y_{t_{j+1}}-y_{t_{j}}\right\rangle_{\mathcal{H}_{\mathcal{S}}^{k-1}(V)}
$$

Hence, computing $k_{\mathcal{S}}(x, y)$ consists in forming the $(P-1) \times(Q-1)$ matrix $M$ such that,

$$
[M]_{i, j}=\left\langle x_{s_{i+1}}-x_{s_{i}}, y_{t_{j+1}}-y_{t_{j}}\right\rangle_{\mathcal{H}_{\mathcal{S}}^{k-1}(V)}
$$

and iteratively applying the update rule as outlined in Alg. 1. Besides, in Alg. 1 we distinguish the case where the solution $u$ on the entire grid is returned, and the case where only the solution at the final times $\left(s_{P}, t_{Q}\right)=(T, T)$ is returned, which corresponds to the value of the kernel $k_{\mathcal{S}}(x, y)$. The runtime complexity to solve one PDE is $\mathcal{O}(P Q)$. We make use of parallelization strategy to drastically speed-up the PDE solver on CUDA-enabled GPUs.

## B. 1 Algorithm for the $1^{\text {st }}$ order MMD

In this section we provide the algorithm to compute an empirical estimator of the $1^{\text {st }}$ order MMD. This way, we introduce subroutines (Alg. 1 and Alg. 2) for the estimator of the $2^{\text {nd }}$ order MMD. We assume that $m=n$ and $P=Q$ to simplify the final runtime complexities of the algorithms.
Let $\left\{x^{i}\right\}_{i=1}^{m} \sim X$ and $\left\{y^{j}\right\}_{j=1}^{n} \sim Y$. An unbiased estimator of the $1^{\text {st }}$ order MMD [1] reads as,

$$
\widehat{\mathcal{D}}_{\mathcal{S}}^{1}(X, Y)=\frac{1}{m(m-1)} \sum_{\substack{i, j=1 \\ i \neq j}}^{m} k_{\mathcal{S}}\left(x^{i}, x^{j}\right)-\frac{2}{m n} \sum_{i, j=1}^{m, n} k_{\mathcal{S}}\left(x^{i}, y^{j}\right)+\frac{1}{n(n-1)} \sum_{\substack{i, j=1 \\ i \neq j}}^{n} k_{\mathcal{S}}\left(y^{i}, y^{j}\right)
$$

Hence, in order to compute this estimator, we need to form the following three Gram matrices $\mathbf{G}_{X, X}^{1} \in \mathbb{R}^{m \times m}, \mathbf{G}_{X, Y}^{1} \in \mathbb{R}^{m \times n}$ and $\mathbf{G}_{Y, Y}^{1} \in \mathbb{R}^{n \times n}$ such that,

$$
\left[\mathbf{G}_{X, X}^{1}\right]_{i, j}=k_{\mathcal{S}}\left(x^{i}, x^{j}\right) \quad\left[\mathbf{G}_{X, Y}^{1}\right]_{i, j}=k_{\mathcal{S}}\left(x^{i}, y^{j}\right) \quad\left[\mathbf{G}_{Y, Y}^{1}\right]_{i, j}=k_{\mathcal{S}}\left(y^{i}, y^{j}\right)
$$

As explained at the begining of this section (and outlined in Alg. 2), this consists in two steps. Taking $\mathbf{G}_{X, Y}^{1}$ for example, first one forms $m \times n$ matrices of size $(P-1) \times(Q-1)$ each of the form,

$$
[M]_{p, q}=\left\langle x_{s_{p+1}}^{i}-x_{s_{p}}^{i}, y_{t_{q+1}}^{j}-y_{t_{q}}^{j}\right\rangle_{V}
$$

and then one solves $m \times n$ PDEs with Alg. 1. The full procedure is summarized in Alg. 3, which has time complexity $\mathcal{O}\left(d m^{2} P^{2}\right)$ where $d$ is the number of coordinates of the paths $x$ and $y$.

```
Algorithm 1 PDESolve \(\mathcal{O}\left(P^{2}\right)\)
    Input: matrix \(M \in \mathbb{R}^{P \times Q}\), full \(\in\{\) True, False \(\}\)
    Output: full solution \(u \in \mathbb{R}^{P \times Q}\) with \(u[p, q]=k_{\mathcal{S}}\left(x_{\left[0, s_{p}\right]}, y_{\left[0, t_{q}\right]}\right)\) or \(u[-1,-1]=k_{\mathcal{S}}(x, y)\)
    \(u[1,:] \leftarrow 1\)
    \(u[:, 1] \leftarrow 1\)
    for \(p\) from 1 to \(P-1\) do
        for \(q\) from 1 to \(Q-1\) do
            \(u[p+1, q+1] \leftarrow f(u[p, q+1], u[p+1, q], u[p, q], M[p, q])\)
    if full then return \(u\) else return \(u[-1,-1]\)
```

```
Algorithm 2 FirstOrderGram \(\mathcal{O}\left(d m^{2} P^{2}\right)\)
    Input: sample paths \(\left\{x^{i}\right\}_{i=1}^{m} \sim X\) and \(\left\{y^{j}\right\}_{j=1}^{n} \sim Y\), full \(\in\{\) True, False \(\}\)
    Output: \(G \in \mathbb{R}^{m \times n \times P \times Q}\) where \(G[i, j, p, q]=k_{\mathcal{S}}\left(x_{\left[0, s_{p}\right]}^{i}, y_{\left[0, t_{q}\right]}^{j}\right)\) or \(G[:,:,-1,-1]\)
    \(M[i, j, p, q] \leftarrow\left\langle x_{s_{p}}^{i}, y_{t_{q}}^{j}\right\rangle \quad \forall i \in\{1, \ldots, m\}, j \in\{1, \ldots, n\}, p \in\{1, \ldots, P\}, q \in\{1, \ldots, Q\}\)
    \(M \leftarrow M[:,:, 1:, 1:]+M[:,:,:-1,:-1]-M[:,:, 1:,:-1]-M[:,:,:-1,1:]\)
    \(G[i, j] \leftarrow \operatorname{PDESolve}(M[i, j]), \quad \forall i \in\{1, \ldots, m\}, j \in\{1, \ldots, n\}\)
    6: if full then return \(G\) else return \(G[:,:,-1,-1]\)
```


## B. 2 Algorithm for the $2^{\text {nd }}$ order MMD

In the main paper, we derived the following estimator of the $2^{\text {nd }}$ order MMD,

$$
\widehat{\mathcal{D}}_{\mathcal{S}}^{2}(X, Y)=\frac{1}{m(m-1)} \sum_{\substack{i, j=1 \\ i \neq j}}^{m} k_{\mathcal{S}}\left(\widetilde{x}^{i}, \widetilde{x}^{j}\right)-\frac{2}{m n} \sum_{i, j=1}^{m, n} k_{\mathcal{S}}\left(\widetilde{x}^{i}, \widetilde{y}^{j}\right)+\frac{1}{n(n-1)} \sum_{\substack{i, j=1 \\ i \neq j}}^{n} k_{\mathcal{S}}\left(\widetilde{y}^{i}, \widetilde{y}^{j}\right)
$$

Compared to the $1^{\text {st }}$ order MMD, in order to compute this estimator, as outlined in Alg. 8, we need to form the following three Gram matrices $\mathbf{G}_{X, X}^{2} \in \mathbb{R}^{m \times m}, \mathbf{G}_{X, Y}^{2} \in \mathbb{R}^{m \times n}$ and $\mathbf{G}_{Y, Y}^{2} \in \mathbb{R}^{n \times n}$,

$$
\left[\mathbf{G}_{X, X}^{2}\right]_{i, j}=k_{\mathcal{S}}\left(\widetilde{x}^{i}, \widetilde{x}^{j}\right) \quad\left[\mathbf{G}_{X, Y}^{2}\right]_{i, j}=k_{\mathcal{S}}\left(\widetilde{x}^{i}, \widetilde{y}^{j}\right) \quad\left[\mathbf{G}_{Y, Y}^{2}\right]_{i, j}=k_{\mathcal{S}}\left(\widetilde{y}^{i}, \widetilde{y}^{j}\right)
$$

As outlined in Alg. 7, this consists in two steps. Taking $\mathbf{G}_{X, Y}^{2}$ for example, first one forms $m \times n$ matrices of size $(P-1) \times(Q-1)$ each of the form,

$$
[M]_{p, q}=\left\langle\widetilde{x}_{s_{p+1}}-\widetilde{x}_{s_{p}}, \widetilde{y}_{t_{q+1}}-\widetilde{y}_{t_{q}}\right\rangle_{\mathcal{H}_{\mathcal{S}}(V)}
$$

(see Alg. 6) and then one solves $m \times n$ PDEs with Alg. 1. This is summarized in Alg. 8, which has time complexity $\mathcal{O}\left((d+m) m^{2} P^{2}\right)$ where $d$ is the number of coordinates of the paths $x$ and $y$.

```
Algorithm 3 FirstOrderMMD
\(\mathcal{O}\left(d m^{2} P^{2}\right)\)
    Input: sample paths \(\left\{x^{i}\right\}_{i=1}^{m} \sim X\) and \(\left\{y^{j}\right\}_{j=1}^{n} \sim Y\)
    Output: an empirical estimator of the \(1^{\text {st }}\) order MMD between \(X\) and \(Y\)
    \(G_{X X}^{1} \leftarrow\) FirstOrderGram \(\left(\left\{x^{i}\right\}_{i=1}^{m},\left\{x^{i}\right\}_{i=1}^{m}\right.\), full = False \()\)
    \(G_{X Y}^{1} \leftarrow\) FirstOrderGram \(\left(\left\{x^{i}\right\}_{i=1}^{m},\left\{y^{j}\right\}_{j=1}^{n}\right.\), full \(=\) False \()\)
    \(G_{Y Y}^{1} \leftarrow\) FirstOrderGram \(\left(\left\{y^{j}\right\}_{j=1}^{n},\left\{y^{j}\right\}_{j=1}^{n},\right.\), full \(=\) False \()\)
    return \(\operatorname{avg}\left(G_{X X}^{1}\right)-2 * \operatorname{avg}\left(G_{X Y}^{1}\right)+\operatorname{avg}\left(G_{Y Y}^{1}\right)\)
```

```
Algorithm 4 SecondOrderGram \(\mathcal{O}\left(P^{2} m^{3}\right)\)
    Input: \(G_{X X}, G_{X Y}, G_{Y Y}\) with \(G_{X Y}[i, j, p, q]=k_{\mathcal{S}}\left(x_{\left[0, s_{p}\right]}^{i}, y_{\left[0, t_{q}\right]}^{j}\right)\) and hyperparameter \(\lambda\).
    Output: an empirical estimator of \(G_{X, Y}^{2} \in \mathbb{R}^{m \times n}\), where \(G_{X, Y}^{2}[i, j]=k_{\mathcal{S}}\left(\widetilde{x}^{i}, \widetilde{y}^{j}\right)\)
    \(M \leftarrow\) InnerProdPredCondKME \(\left(G_{X X}, G_{X Y}, G_{Y Y}, \lambda\right)\)
    \(M \leftarrow M[:,:, 1:, 1:]+M[:,:,:-1,:-1]-M[:,:, 1:,:-1]-M[:,:,:-1,1:]\)
    \(G_{X Y}^{2}[i, j] \leftarrow \operatorname{PDESolve}(M[i, j]), \quad \forall i \in\{1, \ldots, m\}, \forall j \in\{1, \ldots, n\}\)
    return \(G_{X Y}^{2}\)
```

```
Algorithm 5 SecondOrderMMD \(\mathcal{O}\left(d m^{2} P^{2}+P^{2} m^{3}\right)\)
    Input: sample paths \(\left\{x^{i}\right\}_{i=1}^{m} \sim X\) and \(\left\{y^{j}\right\}_{j=1}^{n} \sim Y\), hyperparameter \(\lambda\).
    Output: an empirical estimator of the \(2^{\text {nd }}\) order MMD between \(X\) and \(Y\)
    \(G_{X X}^{1} \leftarrow\) FirstOrderGram \(\left(\left\{x^{i}\right\}_{i=1}^{m},\left\{x^{i}\right\}_{i=1}^{m}\right.\), full \(=\) True \()\)
    \(G_{X Y}^{1} \leftarrow\) FirstOrderGram \(\left(\left\{x^{i}\right\}_{i=1}^{m},\left\{y^{j}\right\}_{j=1}^{n}\right.\), full \(=\) True \()\)
    \(G_{Y Y}^{1} \leftarrow\) FirstOrderGram \(\left(\left\{y^{j}\right\}_{j=1}^{n},\left\{y^{j}\right\}_{j=1}^{n}\right.\), full \(=\) True \()\)
    \(G_{X X}^{2} \leftarrow\) SecondOrderGram \(\left(G_{X X}^{1}, G_{X X}^{1}, G_{X X}^{1}, \lambda\right)\)
    \(G_{X Y}^{2} \leftarrow\) SecondOrderGram \(\left(G_{X X}^{1}, G_{X Y}^{1}, G_{Y Y}^{1}, \lambda\right)\)
    \(G_{Y Y}^{2} \leftarrow \operatorname{SecondOrderGram}\left(G_{Y Y}^{1}, G_{Y Y}^{1}, G_{Y Y}^{1}, \lambda\right)\)
    return \(\operatorname{avg}\left(G_{X X}^{2}\right)-2 * \operatorname{avg}\left(G_{X Y}^{2}\right)+\operatorname{avg}\left(G_{Y Y}^{2}\right)\)
```

```
Algorithm 6 InnerProdPredCondKME \(\mathcal{O}\left(P^{2} m^{3}\right)\)
    Input: three Gram matrices \(G_{X X}, G_{X Y}, G_{Y Y}\) and hyperparameter \(\lambda\)
    Output: returns an empirical estimator of \(M \in \mathbb{R}^{m \times n \times P \times Q}\) where \(M[i, j, p, q]=\left\langle\widetilde{x}_{s_{p}}^{i}, \widetilde{y}_{t_{q}}^{j}\right\rangle\)
    \(W_{X}[:,:, p] \leftarrow\left(G_{X X}[:,:, p, p]+m \lambda I\right)^{-1}, \quad \forall p \in\{1, \ldots, P\}\)
    \(W_{Y}[:,:, q] \leftarrow\left(G_{Y Y}[:,:, q, q]+n \lambda I\right)^{-1}, \forall q \in\{1, \ldots, Q\}\)
    for \(p\) from 1 to \(P\) do
        for \(q\) from 1 to \(Q\) do
            \(M[:,:, p, q] \leftarrow G_{X X}[:,:, p, p]^{T} W_{X}[:,:, p] G_{X Y}[:,:,-1,-1] W_{Y}[:,:, q] G_{Y Y}[:,:, q, q]\)
```


## B. 3 Algorithm for higher order MMDs

Now, we generalize the procedure in Appendix B. 2 for computing an estimator of $\mathcal{D}_{\mathcal{S}}^{k+1}$ when $k>1$,

$$
\widehat{\mathcal{D}}_{\mathcal{S}}^{k+1}(X, Y)=\frac{1}{m(m-1)} \sum_{\substack{i, j=1 \\ i \neq j}}^{m} k_{\mathcal{S}}\left(\widetilde{x}^{k, i}, \widetilde{x}^{k, j}\right)-\frac{2}{m n} \sum_{i, j=1}^{m, n} k_{\mathcal{S}}\left(\widetilde{x}^{k, i}, \widetilde{y}^{k, j}\right)+\frac{1}{n(n-1)} \sum_{\substack{i, j=1 \\ i \neq j}}^{n} k_{\mathcal{S}}\left(\widetilde{y}^{k, i}, \widetilde{y}^{k, j}\right),
$$

where $\widetilde{x}^{k, i}$ and $\widetilde{y}^{k, j}$ denote sample paths from the processes $\mu_{X \mid \mathcal{F}_{X}}^{k}$ and $\mu_{Y \mid \mathcal{F}_{Y}}^{k}$ respectively. In order to compute this estimator, as outlined in Alg. 8, we need to form the following three Gram matrices $\mathbf{G}_{X, X}^{k+1} \in \mathbb{R}^{m \times m}, \mathbf{G}_{X, Y}^{k+1} \in \mathbb{R}^{m \times n}$ and $\mathbf{G}_{Y, Y}^{k+1} \in \mathbb{R}^{n \times n}$,

$$
\left[\mathbf{G}_{X, X}^{k+1}\right]_{i, j}=k_{\mathcal{S}}\left(\widetilde{x}^{k, i}, \widetilde{x}^{k, j}\right) \quad\left[\mathbf{G}_{X, Y}^{k+1}\right]_{i, j}=k_{\mathcal{S}}\left(\widetilde{x}^{k, i}, \widetilde{y}^{k, j}\right) \quad\left[\mathbf{G}_{Y, Y}^{k+1}\right]_{i, j}=k_{\mathcal{S}}\left(\widetilde{y}^{k, i}, \widetilde{y}^{k, j}\right)
$$

As outlined in Alg. 7, this consists in two steps. Taking $\mathbf{G}_{X, Y}^{k+1}$ for example, first one forms $m \times n$ matrices of size $(P-1) \times(Q-1)$ each of the form,

$$
[M]_{p, q}=\left\langle\widetilde{x}_{s_{p+1}}^{k}-\widetilde{x}_{s_{p}}^{k}, \widetilde{y}_{t_{q+1}}^{k}-\widetilde{y}_{t_{q}}^{k}\right\rangle_{\mathcal{H}_{\mathcal{S}}^{k}(V)}
$$

(see Alg. 6) and then one solves $m \times n$ PDEs with Alg. 1. This is summarized in Alg. 8, which has time complexity $\mathcal{O}\left((d+k m) m^{2} P^{2}\right)$ where $d$ is the number of coordinates of the paths $x$ and $y$.

```
Algorithm 7 HigherOrderGram \(\mathcal{O}\left(P^{2} m^{3}\right)\)
    Input: \(G_{X X}^{k}, G_{X Y}^{k}, G_{Y Y}^{k}\) with \(G_{X Y}^{k}[i, j, p, q]=k_{\mathcal{S}}\left(\widetilde{x}_{\left[0, s_{p}\right]}^{k-1, i}, \widetilde{y}_{\left[0, t_{q}\right]}^{k-1, j}\right)\) and hyperparameter \(\lambda\).
    Output: an estimator of \(G_{X, Y}^{k+1} \in \mathbb{R}^{m \times n \times P \times Q}\), where \(G_{X, Y}^{k+1}[i, j, p, q]=k_{\mathcal{S}}\left(\widetilde{x}_{\left[0, s_{p}\right]}^{k, i}, \widetilde{y}_{\left[0, t_{q}\right]}^{k, j}\right)\)
    \(M \leftarrow\) InnerProdPredCondKME \(\left(G_{X X}^{k}, G_{X Y}^{k}, G_{Y Y}^{k}, \lambda\right)\)
    \(M \leftarrow M[:,:, 1:, 1:]+M[:,:,:-1,:-1]-M[:,:, 1:,:-1]-M[:,:,:-1,1:]\)
    \(G_{X Y}^{k+1}[i, j] \leftarrow \operatorname{PDESolve}(M[i, j]\), full \(=\) True \(), \quad \forall i \in\{1, \ldots, m\}, \forall j \in\{1, \ldots, n\}\)
    return \(G_{X Y}^{k+1}\)
```

```
Algorithm 8 HigherOrderMMD \(\quad \mathcal{O}\left(d m^{2} P^{2}+(k-1) P^{2} m^{3}\right)\)
    Input: sample paths \(\left\{x^{i}\right\}_{i=1}^{m} \sim X\) and \(\left\{y^{j}\right\}_{j=1}^{n} \sim Y\), hyperparameter \(\lambda\), order \(k\).
    Output: an empirical estimator of the \(k^{\text {th }}\) order MMD between \(X\) and \(Y\)
    \(G_{X X} \leftarrow\) FirstOrderGram \(\left(\left\{x^{i}\right\}_{i=1}^{m},\left\{x^{i}\right\}_{i=1}^{m}\right.\), full \(=\) True \()\)
    \(G_{X Y} \leftarrow\) FirstOrderGram \(\left(\left\{x^{i}\right\}_{i=1}^{m},\left\{y^{j}\right\}_{j=1}^{n}\right.\), full = True \()\)
    \(G_{Y Y} \leftarrow\) FirstOrderGram \(\left(\left\{y^{j}\right\}_{j=1}^{n},\left\{y^{j}\right\}_{j=1}^{n}\right.\), full \(=\) True \()\)
    for order from 2 to \(k\) do
        \(G_{X X}^{\text {new }} \leftarrow\) HigherOrderGram \(\left(G_{X X}, G_{X X}, G_{X X}, \lambda\right)\)
        \(G_{X Y}^{\text {new }} \leftarrow\) HigherOrderGram \(\left(G_{X X}, G_{X Y}, G_{Y Y}, \lambda\right)\)
        \(G_{Y Y}^{\text {new }} \leftarrow\) HigherOrderGram \(\left(G_{Y Y}, G_{Y Y}, G_{Y Y}, \lambda\right)\)
        \(G_{X X}, G_{X Y}, G_{Y Y} \leftarrow G_{X X}^{\text {new }}, G_{X Y}^{\text {new }}, G_{Y Y}^{\text {new }}\)
    return \(\operatorname{avg}\left(G_{X X}[:,:,-1,-1]\right)-2 * \operatorname{avg}\left(G_{X Y}[:,:,-1,-1]\right)+\operatorname{avg}\left(G_{Y Y}[:,:,-1,-1]\right)\)
```


## C Experimental details

We start with further experimental details for the applications of higher order distribution regression to quantitative finance (Sec. 4.2), where we consider the problem of optimally stopping fractional Brownian motions with different hurst exponents.

## C. 1 Rough volatility

Rough volatility models constitute a class of models that are empirically well-tailored to fit observed implied market volatilities in the context of option pricing for short maturity assets. The basic model for option pricingis called the Black-Scholes model in which the volatility is assumed to be constant. Stochastic volatility models are extensions of the Black-Scholes model to the case where the volatility is itself stochastic. The main shortcoming of such stochastic volatility models is that they are able to capture the true steepness of the implied volatility smile close to maturity (see [36] for extra details). This is where rough volatility models become useful. Among them, the rough Bergomi model introduced by [36], stood out for its ability to explain implied volatility and other phenomena related to European options.

## C. 2 Higher order distribution regression

Data We use the data generator from https://github.com/HeKrRuTe/OptStopRandNN to simulate sample paths from $X$ a fractional Brownian motion ( fBm ) and obtain the solution of the optimal stopping time problem $\sup _{\tau} \mathbb{E}\left[g\left(X_{\tau}\right) \mid X_{0}\right]$. We note that although fBm is not typically used as a stock price model in quantitative finance, it is nevertheless considered a respected challenging example for optimal stopping algorithms [51, 52].

Models We use a kernel Ridge regressor with different distribution regression kernels. Each is of the form $K(X, Y)=\exp \left(-\mathcal{D}(X, Y)^{2} / \sigma^{2}\right)$ where $\mathcal{D}(X, Y)$ is a maximum mean discrepancy. The models $K_{\mathcal{S}}^{1}$ and $K_{\mathcal{S}}^{2}$ correspond to the $1^{\text {st }}$ and $2^{\text {nd }}$ order maximum mean discrepancies $\mathcal{D}_{\mathcal{S}}^{1}$ and $\mathcal{D}_{\mathcal{S}}^{2}$. We consider two other baselines (Matérn and RBF) for which the MMD is computed using the

Matern 3/2 covariance function $k_{\text {mat32 }}$, and the RBF covariance function $k_{\mathrm{rbf}}$,

$$
k_{\mathrm{mat} 32}(x, y)=\left(1+\frac{\sqrt{3}}{\gamma^{2}}\|x-y\|\right) \exp \left(-\frac{\sqrt{3}}{\gamma^{2}}\|x-y\|\right), \quad k_{\mathrm{rbf}}(x, y)=\exp \left(-\frac{\|x-y\|^{2}}{\gamma^{2}}\right)
$$

All models are run 3 times. The hyperparameters of all models are selected by cross-validation via a grid search on the training set ( $70 \%$ of the data selected at random) of each run.

## C.2.1 Inferring causal graph for interacting bodies

We provide further details for the last application (Sec. 4.3) where the task is to infer whether any two bodies are connected by a spring from multiple observations of their 2D trajectories.

Data We adapt the Pymunk simulator from [42] publicly available at https://github.com/ pairlab/v-cdn. For each pair of balls, there is a one-half probability that they are connected by nothing, or a spring. For each graph we run multiple episodes each of 20 time steps. At the beginning of each episode, we randomly assign the balls in different positions. The stiffness of the spring relation is set to 20 , and we randomly sample the rest length between [20, 120].

Causal discovery algorithm The PC algorithm [53] uses conditional independence tests to generate a causal graph from a dataset. The PC algorithm consists in two stages. The first stage, referred to as the skeleton phase, consists in finding the structure of the causal graph. In the second stage, the edges are oriented by repetitively applying orientation rules. In the multi-body interaction example, we only need to perform the skeleton phase, which is sketched hereafter,

1. Start with a complete graph
2. For each $X$ and $Y$ which are still connected. If there is a third variable $Z_{1}$ connected to $X$ or $Y$, such that $X \Perp Y \mid Z_{1}$, remove the edge between $X$ and $Y$.
3. For each $X$ and $Y$ which are still connected, if there is a third and a fourth variable $Z_{1}$ and $Z_{2}$ connected to $X$ or $Y$ such that $X \Perp Y \mid Z_{1}, Z_{2}$, remove the edge between $X$ and $Y$.
4. Iteratively increase the cardinality of the set of variables on which to condition.

To test for conditional independence we use the Hilbert-Schmidt conditional independence criterion $H_{X Y \mid Z}$ for stochastic processes (Appendix A.1). The combination of the PC algorithm with a kernel-based dependence measure has been used in [13] and [12] where it is termed kPC. However, to our knowledge it has never been used in conjunction with a kernel-based measure of dependence for multidimensional stochastic processes.
Since the null distribution of the test statistics $H_{Y X \mid Z}$ is not known, one possibility would be to use a permutation approach as in [12, Sec 2]. However the latter is not computationally efficient. We leave the development of a faster approach for future work, and adopt the approach [13] for this experiment. That is we use a threshold $\alpha$ and remove an edge if there is a $Z$ such that $H_{X Y \mid Z}<\alpha$. We repeat 15 times the causal discovery procedure and use $30 \%$ of the runs to fix $\alpha$.
All experiments in Sec. 4 have been run on a P100 GPU to leverage an efficient dedicated CUDA implementation of the signature kernel.

## D Proofs

Theorem 6. Given two stochastic processes $X, Y$

$$
\mathcal{D}_{\mathcal{S}}^{2}(X, Y)=0 \Longleftrightarrow \mathbb{P}_{X \mid \mathcal{F}_{X}}=\mathbb{P}_{Y \mid \mathcal{F}_{Y}}
$$

Furthermore

$$
\mathcal{D}_{\mathcal{S}}^{2}(X, Y)=0 \Longrightarrow \mathcal{D}_{\mathcal{S}}^{1}(X, Y)=0
$$

but the converse is not generally true.
Proof. First we note that by a standard result in signature kernel learning theory. e.g., [31], for $X \in \mathcal{X}(V)$ and every $t$, the mapping

$$
\mathbb{P}_{X \mid \mathcal{F}_{X_{t}}} \mapsto \mu_{X \mid \mathcal{F}_{X_{t}}}^{1}=\int k_{\mathcal{S}}(\cdot, x) \mathbb{P}_{X \mid \mathcal{F}_{X_{t}}}(d x)
$$

is a homeomorphism (with respect to weak topology and Hilbert space topology); in particular, we have

$$
\mathbb{P}_{X \mid \mathcal{F}_{X}}=\mathbb{P}_{Y \mid \mathcal{F}_{Y}} \Longleftrightarrow \mathbb{P}_{\mu_{X \mid F_{X}}^{1}}=\mathbb{P}_{\mu_{Y \mid \mathcal{F}_{Y}}^{1}} .
$$

Then, using the same argument for $\mathbb{P}_{\mu_{X \mid \mathcal{F}_{X}}^{1}}$ and $\mathbb{P}_{\mu_{Y \mid \mathcal{F}_{Y}}^{1}}$, we can further deduce that

$$
\mathbb{P}_{\mu_{X \mid \mathcal{F}_{X}}^{1}}=\mathbb{P}_{\mu_{Y \mid \mathcal{F}_{Y}}^{1}} \Longleftrightarrow \int k_{\mathcal{S}}(\cdot, x) \mathbb{P}_{\mu_{X \mid \mathcal{F}_{X}}^{1}}(d x)\left(=\mu_{X}^{2}\right)=\int k_{\mathcal{S}}(\cdot, y) \mathbb{P}_{\mu_{Y \mid \mathcal{F}_{Y}}^{1}}(d y)\left(=\mu_{Y}^{2}\right)
$$

Since by definition it holds that $\mathcal{D}_{\mathcal{S}}^{2}(X, Y)=\left\|\mu_{X}^{2}-\mu_{Y}^{2}\right\|_{\mathcal{H}^{2}(V)}$, we complete the proof of the first claim in this theorem.

For the second claim, it is easy to see that by definition $\mathbb{P}_{X \mid \mathcal{F}_{X}}=\mathbb{P}_{Y \mid \mathcal{F}_{Y}}$ ensures that $\mathbb{P}_{X}=\mathbb{P}_{Y}$, and therefore the implication that $\mathcal{D}_{\mathcal{S}}^{2}(X, Y)=0 \Longrightarrow \mathcal{D}_{\mathcal{S}}^{1}(X, Y)=0$ follows immediately from the fact that $\mathbb{P}_{X}=\mathbb{P}_{Y} \Longleftrightarrow \mathcal{D}_{\mathcal{S}}^{1}(X, Y)=0$. Moreover, we refer readers to [26, Example 3.1] for a simple example which shows that there exist processes $X$ and $Y$ with $\mathcal{D}_{\mathcal{S}}^{1}(X, Y)=0$ but $\mathcal{D}_{\mathcal{S}}^{2}(X, Y)>0$.

Theorem 7. $\widehat{\mathcal{D}}_{\mathcal{S}}^{2}(X, Y)$ is a consistent estimator for the $2^{\text {nd }}$ order MMD, i.e.

$$
\begin{equation*}
\left|\widehat{\mathcal{D}}_{\mathcal{S}}^{2}(X, Y)-\mathcal{D}_{\mathcal{S}}^{2}(X, Y)\right| \xrightarrow{p} 0 \quad \text { as } m, n \rightarrow \infty \tag{39}
\end{equation*}
$$

with $\left\{x^{i}\right\}_{i=1}^{m} \sim X,\left\{y^{i}\right\}_{i=1}^{n} \sim Y$ and where convergence is in probability.
Proof. Recall that given $m$ independent sample paths $\left\{x^{i}\right\}_{i=1}^{m} \sim X$, we can use the estimator in appendix A. 2 to approximate sample paths $\left\{\widetilde{x}^{i}\right\}_{i=1}^{m}$ from the $1^{\text {st }}$ order predictive KME $\mu_{X \mid \mathcal{F}_{X}}^{1}$. Hence, it suffices to prove the following claim.

Claim: consider $m$ independent sample paths $\left\{\widetilde{x}^{i}\right\}_{i=1}^{m} \sim \mu_{X \mid \mathcal{F}_{X}}^{1}$. Then, the estimator given by $\widehat{\mu}_{X}^{2}=\frac{1}{m} \sum_{i=1}^{m} k_{\mathcal{S}}\left(\widetilde{x}^{i}, \cdot\right)$ is consistent for the $2^{\text {nd }}$ order predictive KME $\mu_{X}^{2}$, i.e.

$$
\begin{equation*}
\left\|\widehat{\mu}_{X}^{2}-\mu_{X}^{2}\right\|_{\mathcal{H}_{\mathcal{S}}^{2}(V)}^{2} \xrightarrow{p} 0, \text { as } m \rightarrow \infty \tag{40}
\end{equation*}
$$

By the triangular inequality

$$
\begin{align*}
\left\|\widehat{\mu}_{X}^{2}-\mu_{X}^{2}\right\|_{\mathcal{H}_{\mathcal{S}}^{2}(V)}^{2} & =\left\|\frac{1}{m} \sum_{i=1}^{m} k_{\mathcal{S}}\left(\widetilde{x}^{i}, \cdot\right)-\mathbb{E}\left[k_{\mathcal{S}}\left(\mu_{X \mid \mathcal{F}_{X}}^{1}, \cdot\right)\right]\right\|_{\mathcal{H}_{\mathcal{S}}^{2}(V)}^{2}  \tag{41}\\
& \leq\left\|\frac{1}{m} \sum_{i=1}^{m} k_{\mathcal{S}}\left(\widetilde{x}^{i}, \cdot\right)-\mathbb{E}\left[k_{\mathcal{S}}\left(\widehat{\mu}_{X \mid \mathcal{F}_{X}}^{1}, \cdot\right)\right]\right\|_{\mathcal{H}_{\mathcal{S}}^{2}(V)}^{2}  \tag{42}\\
& +\left\|\mathbb{E}\left[k_{\mathcal{S}}\left(\widehat{\mu}_{X \mid \mathcal{F}_{X}}^{1}, \cdot\right)\right]-\mathbb{E}\left[k_{\mathcal{S}}\left(\mu_{X \mid \mathcal{F}_{X}}^{1}, \cdot\right)\right]\right\|_{\mathcal{H}_{\mathcal{S}}^{2}(V)}^{2} \tag{43}
\end{align*}
$$

The term in (42) converges to 0 as $m \rightarrow \infty$ by the weak law of large numbers. Therefore, it remains to show that

$$
\begin{equation*}
\left\|\mathbb{E}_{X}\left[k_{\mathcal{S}}\left(\widehat{\mu}_{X \mid \mathcal{F}_{X}}^{1}, \cdot\right)\right]-\mathbb{E}_{X}\left[k_{\mathcal{S}}\left(\mu_{X \mid \mathcal{F}_{X}}^{1}, \cdot\right)\right]\right\|_{\mathcal{H}_{\mathcal{S}}^{2}(V)} \xrightarrow{p} 0, \text { as } m \rightarrow \infty \tag{44}
\end{equation*}
$$

First note the following upper bound

$$
\left\|\mathbb{E}_{X}\left[k_{\mathcal{S}}\left(\widehat{\mu}_{X \mid \mathcal{F}_{X}}^{1}, \cdot\right)\right]-\mathbb{E}_{X}\left[k_{\mathcal{S}}\left(\mu_{X \mid \mathcal{F}_{X}}^{1}, \cdot\right)\right]\right\|_{\mathcal{H}_{\mathcal{S}}^{2}(V)} \leq \mathbb{E}_{X}\left\|k_{\mathcal{S}}\left(\widehat{\mu}_{X \mid \mathcal{F}_{X}}^{1}, \cdot\right)-k_{\mathcal{S}}\left(\mu_{X \mid \mathcal{F}_{X}}^{1}, \cdot\right)\right\|_{\mathcal{H}_{\mathcal{S}}^{2}(V)}
$$

We will show convergence of the right-hand-side. By [54, Theorem 3.4], for every $t=1, \ldots, T$

$$
\begin{equation*}
\mathbb{E}_{X \mid \mathcal{F}_{X_{t}}}\left\|\widehat{\mu}_{X \mid \mathcal{F}_{X_{t}}}^{1}-\mu_{X \mid \mathcal{F}_{X_{t}}}^{1}\right\|_{\mathcal{H}_{\mathcal{S}}(V)}^{2} \xrightarrow{p} 0 \quad \text { as } m \rightarrow \infty \tag{45}
\end{equation*}
$$

Now let us assume that the above convergences also hold almost surely for every $t=1, \ldots, T$. Then by the Egorov's theorem, for any $\delta>0$, there is a subset $\Omega_{\delta}$ with $\mathbb{P}\left(\Omega_{\delta}\right)>1-\delta$ and the above
convergence (45) holds uniformly on $\Omega_{\delta}$. This implies that on $\Omega_{\delta}$ for every $\varepsilon>0$ there is an $N(\varepsilon)$ such that for all $m \geq N(\varepsilon)$ and all $t=1, \ldots, T$, it holds that

$$
\begin{equation*}
\mathbb{E}_{X \mid \mathcal{F}_{X_{t}}}\left\|\widehat{\mu}_{X \mid \mathcal{F}_{X_{t}}}^{1}-\mu_{X \mid \mathcal{F}_{X_{t}}}^{1}\right\|_{\mathcal{H}_{\mathcal{S}}(V)}^{2} \leq \varepsilon^{2} \tag{46}
\end{equation*}
$$

From this estimate we immediately obtain by the triangle inequality on $\Omega_{\delta}$

$$
\begin{equation*}
\mathbb{E}_{X \mid \mathcal{F}_{X_{t}}}\left\|\widehat{\mu}_{X \mid \mathcal{F}_{X_{t}}}^{1}\right\|_{\mathcal{H}_{s}(V)}^{2} \leq 2 \mathbb{E}_{X \mid \mathcal{F}_{X_{t}}}\left\|\mu_{X \mid \mathcal{F}_{X_{t}}}^{1}\right\|_{\mathcal{H}_{\mathcal{S}}(V)}^{2}+2 \varepsilon^{2} \tag{47}
\end{equation*}
$$

and, by the Chebyshev's inequality, on $\Omega_{\delta}, \forall t=1, \ldots, T, \forall m \geq N(\varepsilon)$
$\mathbb{P}_{X \mid \mathcal{F}_{X_{t}}}\left[\left\|\widehat{\mu}_{X \mid \mathcal{F}_{X_{t}}}^{1}-\mu_{X \mid \mathcal{F}_{X_{t}}}^{1}\right\|_{\mathcal{H}_{\mathcal{S}}(V)}>\sqrt{\varepsilon}\right] \leq \frac{1}{\varepsilon} \mathbb{E}_{X \mid \mathcal{F}_{X_{t}}}\left\|\widehat{\mu}_{X \mid \mathcal{F}_{X_{t}}}^{1}-\mu_{X \mid \mathcal{F}_{X_{t}}}^{1}\right\|_{\mathcal{H}_{\mathcal{S}}(V)}^{2} \leq \frac{1}{\varepsilon} \varepsilon^{2}=\varepsilon$,
which implies that on $\Omega_{\delta}, \forall t=1, \ldots, T$, the sequence $\widehat{\mu}_{X \mid \mathcal{F}_{X_{t}}}^{1}$ converges to $\mu_{X \mid \mathcal{F}_{X_{t}}}^{1}$ in probability with respect to $\mathbb{P}_{X}$. By a standard result in rough path theory [30] there exists a universal constant $\beta \in \mathbb{R}$ such that

$$
\begin{equation*}
\left\|k_{\mathcal{S}}\left(\widehat{\mu}_{X \mid \mathcal{F}_{X}}^{1}, \cdot\right)\right\|_{\mathcal{H}_{\mathcal{S}}^{2}(V)} \leq \beta\left\|\widehat{\mu}_{X \mid \mathcal{F}_{X}}^{1}\right\|_{\mathcal{H}_{\mathcal{S}}(V)}^{1-\mathrm{var}} \tag{48}
\end{equation*}
$$

where $\|\cdot\|_{\mathcal{H}_{\mathcal{S}}(V)}^{1-\mathrm{var}}$ denotes the total variation norm of paths taking values in $\mathcal{H}_{\mathcal{S}}(V)$. Since we are in a finite discrete time setup, it is easy to see that

$$
\begin{equation*}
\left\|\widehat{\mu}_{X \mid \mathcal{F}_{X}}^{1}\right\|_{\mathcal{H}_{\mathcal{S}}(V)}^{1-\mathrm{var}} \leq C(T) \sum_{t=1}^{T}\left\|\widehat{\mu}_{X \mid \mathcal{F}_{X_{t}}}^{1}\right\|_{\mathcal{H}_{\mathcal{S}}(V)} \tag{49}
\end{equation*}
$$

Hence, combining all above arguments, we can conclude that on $\Omega_{\delta}$ and for all $m \geq N(\varepsilon)$,

$$
\begin{align*}
\mathbb{E}_{X}\left\|k_{\mathcal{S}}\left(\widehat{\mu}_{X \mid \mathcal{F}_{X}}^{1}, \cdot\right)\right\|_{\mathcal{H}_{\mathcal{S}}^{2}(V)}^{2} & \leq \beta^{2} \mathbb{E}_{X}\left\|\widehat{\mu}_{X \mid \mathcal{F}_{X}}^{1}\right\|_{1-\mathrm{var} ; \mathcal{H}_{\mathcal{S}}(V)}^{2}  \tag{50}\\
& \leq \beta^{2} C(T) \sum_{t=1}^{T} \mathbb{E}_{X}\left\|\widehat{\mu}_{X \mid \mathcal{F}_{X_{t}}}^{1}\right\|_{\mathcal{H}_{\mathcal{S}}(V)}^{2}  \tag{51}\\
& \leq \beta^{2} C(T)\left(2 \sum_{t=1}^{T} \mathbb{E}_{X}\left\|\mu_{X \mid \mathcal{F}_{X_{t}}}^{1}\right\|_{\mathcal{H}_{\mathcal{S}}(V)}^{2}+2 T \varepsilon^{2}\right) \leq C<\infty \tag{52}
\end{align*}
$$

where in the last line we used (47). As a result, we obtain that on $\Omega_{\delta}$,

$$
\begin{equation*}
\sup _{m \geq N(\varepsilon)} \mathbb{E}_{X}\left\|k_{\mathcal{S}}\left(\widehat{\mu}_{X \mid \mathcal{F}_{X}}^{1}, \cdot\right)-k_{\mathcal{S}}\left(\mu_{X \mid \mathcal{F}_{X}}^{1}, \cdot\right)\right\|_{\mathcal{H}_{\mathcal{S}}^{2}(V)}^{2}<\infty \tag{53}
\end{equation*}
$$

which in turn implies, by the de la Vallée-Poussin theorem, that on $\Omega_{\delta}$, the sequence $\left\|k_{\mathcal{S}}\left(\widehat{\mu}_{X \mid \mathcal{F}_{X}}^{1}, \cdot\right)-k_{\mathcal{S}}\left(\mu_{X \mid \mathcal{F}_{X}}^{1}, \cdot\right)\right\|_{\mathcal{H}_{\mathcal{S}}^{2}(V)}^{2}, m \geq N(\varepsilon)$ is uniformly integrable for $\mathbb{P}_{X}$. Then recalling that we have shown that $\widehat{\mu}_{X \mid \mathcal{F}_{X}}^{1}$ converges to $\mu_{X \mid \mathcal{F}_{X}}^{1}$ in probability with respect to $\mathbb{P}_{X}$ as $m \rightarrow \infty$, a standard result in probability theory ensures that (thanks to the uniform integrability of the sequence and the continuity of the kernel $k_{\mathcal{S}}$ which ensures that $\left\|k_{\mathcal{S}}\left(\widehat{\mu}_{X \mid \mathcal{F}_{X}}^{1}, \cdot\right)-k_{\mathcal{S}}\left(\mu_{X \mid \mathcal{F}_{X}}^{1}, \cdot\right)\right\|_{\mathcal{H}_{\mathcal{S}}^{2}(V)} \rightarrow 0$ in probability for $\left.\mathbb{P}_{X}\right)$ on $\Omega_{\delta}, \mathbb{E}_{X}\left\|k_{\mathcal{S}}\left(\widehat{\mu}_{X \mid \mathcal{F}_{X}}^{1}, \cdot\right)-k_{\mathcal{S}}\left(\mu_{X \mid \mathcal{F}_{X}}^{1}, \cdot\right)\right\|_{\mathcal{H}_{\mathcal{S}}^{2}(V)} \rightarrow 0$, as $m \rightarrow \infty$, which implies that $\left\|\mathbb{E}_{X}\left[k_{\mathcal{S}}\left(\widehat{\mu}_{X \mid \mathcal{F}_{X}}^{1}, \cdot\right)\right]-\mathbb{E}_{X}\left[k_{\mathcal{S}}\left(\mu_{X \mid \mathcal{F}_{X}}^{1}, \cdot\right)\right]\right\|_{\mathcal{H}_{\mathcal{S}}^{2}(V)} \rightarrow 0$, as $m \rightarrow \infty$, on $\Omega_{\delta}$. Clearly, as $\delta$ was arbitrary, we have $\mathbb{E}_{X \mid \mathcal{F}_{X_{t}}}\left\|\widehat{\mu}_{X \mid \mathcal{F}_{X_{t}}}^{1}-\mu_{X \mid \mathcal{F}_{X_{t}}}^{1}\right\|_{\mathcal{H}_{s}(V)}^{2} \rightarrow 0 \quad$ as $m \rightarrow \infty$ a.s. Finally, note that the above result holds true for any subsequence of $\left(\widehat{\mu}_{X \mid \mathcal{F}_{X_{t}}}^{1}\right)_{m \geq 1}$ (because every sequence converging in probability has a subsequence converging almost surely), which proves the desired result (40).

Theorem 8. Given two stochastic processes $X, Y$

$$
\begin{equation*}
\mathcal{D}_{\mathcal{S}}^{n}(X, Y)=0 \Longrightarrow \mathcal{D}_{\mathcal{S}}^{k}(X, Y)=0 \quad \text { for any } 1<k<n \tag{54}
\end{equation*}
$$

but the converse is not generally true.
Proof. Let $X \in \mathcal{X}(V)$, then we denote $\mathbb{P}_{X \mid \mathcal{F}_{X_{t}}}=: X_{t}^{(1)}$. Then we continue this procedure and define $X_{t}^{(n)}:=\mathbb{P}_{X^{n-1} \mid \mathcal{F}_{X_{t}}}$ (it is called the rank $n$ prediction process in [29]). Now we can apply the same argument as in the proof of Theorem 6 together with an induction procedure easily prove that

$$
\begin{equation*}
\mathcal{D}_{\mathcal{S}}^{n}(X, Y)=0 \Longleftrightarrow \mathbb{P}_{X^{(n-1)}}=\mathbb{P}_{Y^{(n-1)}} \tag{55}
\end{equation*}
$$

for all $n>1$. From the definition of these processes $X^{(n)}$ and $Y^{(n)}$ we can immediately see that $\mathbb{P}_{X^{(n)}}=\mathbb{P}_{Y^{(n)}}$ ensures that $\mathbb{P}_{X^{(k)}}=\mathbb{P}_{Y^{(k)}}$ for all $k<n$, which yields the desired result. We refer readers to [26, Example 3.2] for examples which illustrate that for each $n$ there exist processes $X$ and $Y$ with $\mathcal{D}_{\mathcal{S}}^{n}(X, Y)=0$ (equivalently, $\mathbb{P}_{X^{(n-1)}}=\mathbb{P}_{Y^{(n-1)}}$ ) but $\mathcal{D}_{\mathcal{S}}^{n+1}(X, Y)>0$ (equivalently, $\left.\mathbb{P}_{X^{(n)}} \neq \mathbb{P}_{Y^{(n)}}\right)$.

Remark. Using terminologies from [26] and [29], $\mathbb{P}_{X^{(n)}}=\mathbb{P}_{Y^{(n)}}$ means that processes $X$ and $Y$ have the same adapted distribution up to rank $n$. Therefore Thm. 6 and 8 tell us that $\mathcal{D}_{\mathcal{S}}^{n}(X, Y)=0$ if and only if they ave the same adapted distribution up to rank $n$. Moreover, using the partial isometry between RKHS generated by $k_{\mathcal{S}}$ and tensor algebra space, see e.g. [31, Theorem E.2], one can use an induction argument to verify that $\mathcal{D}_{\mathcal{S}}^{n}$ coincides with the metric $d_{n-1}$ defined in [29, Definition 14], and therefore by [29, Theorem 4] we can obtain a stronger result that $\mathcal{D}_{\mathcal{S}}^{n}$ actually metrizes the so-called rank $n-1$ adapted topology (see [29, Definition 5], [26, Definition 2.25]). For more details regarding adapted topologies we refer to [29].
Theorem 9. Let $f: \mathbb{R} \rightarrow \mathbb{R}$ be a globally analytic function with non-negative coefficients. Define the family of kernels $K_{\mathcal{S}}^{n}: \mathcal{P}(\mathcal{X}(V)) \times \mathcal{P}(\mathcal{X}(V)) \rightarrow \mathbb{R}$ as follows

$$
\begin{equation*}
K_{\mathcal{S}}^{n}(X, Y)=f\left(\mathcal{D}_{\mathcal{S}}^{n}(X, Y)\right), \quad n \in \mathbb{N}_{\geq 1} \tag{56}
\end{equation*}
$$

Then the RKHS associated to $K_{\mathcal{S}}^{n}$ is dense in the space of functions from $\mathcal{P}(\mathcal{X}(V))$ to $\mathbb{R}$ which are continuous with respect to the $k^{\text {th }}$ order MMD for any $1 \leq k \leq n$.

Proof. By [55, Thm. 2.2] if $K$ is a compact metric space and $H$ is a separable Hilbert space such that there exists a continuous (w.r.t. a topology $\tau$ on $K$ ) and injective map $\rho: K \rightarrow H$, then for any globally analytic function with non-negative coefficients $f: \mathbb{R} \rightarrow \mathbb{R}$ the kernel $k: K \times K \rightarrow \mathbb{R}$ given by

$$
\begin{equation*}
k\left(z, z^{\prime}\right)=f\left(\left\|\rho(z)-\rho\left(z^{\prime}\right)\right\|_{H}\right) \tag{57}
\end{equation*}
$$

is universal in the sense that its RKHS is $\tau$-dense in the space of $\tau$-continuous functions from $K$ to $\mathbb{R}$. By assumption, $\mathcal{X}(V)$ is a $\mathcal{D}_{\mathcal{S}}^{1}$-compact metric space, therefore by Thm. 8 it is also $\mathcal{D}_{\mathcal{S}}^{n}$-compact for every $n \geq 1$. Hence, by [56, Thm. 10.2] the set of stochastic processes $\mathcal{P}(\mathcal{X}(V))$ is also $\mathcal{D}_{\mathcal{S}}^{n}$-compact. For showing that $\rho: X \mapsto \mu_{X}^{n}$ is injective and continous with respect to $\mathcal{D}_{\mathcal{S}}^{n}$ we refer readers to [29, Proposition 4], one only needs to verify that $\mu_{X}^{n}$ corresponds to the mapping $\bar{S}_{n}$ used in the proof of [29, Proposition 4] by the definition of $\mu_{X}^{n}$ and the fact that $\mathcal{D}_{\mathcal{S}}^{n}$ metrizes the rank $n-1$ adapted topology (cf. the above remark). Furthermore $\mathcal{H}_{\mathcal{S}}^{n}(V)$ can be shown by induction to be a Hilbert space with a countable basis, hence it is separable. Setting $K=\mathcal{P}(\mathcal{X}(V)), H=\mathcal{H}_{\mathcal{S}}^{n}(V)$ and $\rho: X \mapsto \mu_{X}^{n}$ concludes the proof.

