Path-Dependent McKean–Vlasov optimal control.

Applications to Time-Series Generation

Samy Mekkaoui

CMAP, École Polytechnique

International Conference on Computational Finance

1–4 September 2026

Oxford, United Kingdom

Missing video from Huyen's talk yesterday

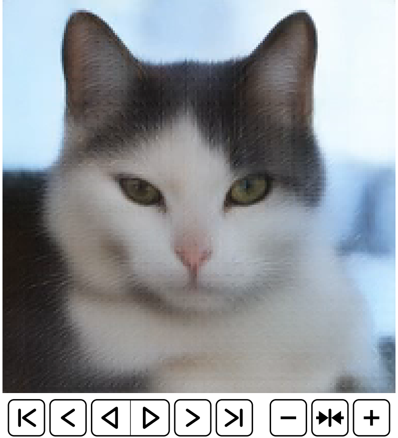

Outline

Introduction and motivation

Path-dependent Mc-Kean Vlasov optimal control

Financial time series generation

Generative modeling through optimal control

Generative modeling can be viewed as an **optimization problem over probability distributions**: one seeks a mechanism whose output law matches the distribution of the observed data.

Generative modeling through optimal control

Generative modeling can be viewed as an **optimization problem over probability distributions**: one seeks a mechanism whose output law matches the distribution of the observed data.

- For time series, matching only marginal distributions is generally not enough: the generator should reproduce temporal, path-dependent, and volatility features.
- Stochastic control provides a natural way to **steer the law of a stochastic process** while retaining an interpretable dynamical model.
- When the objective depends explicitly on the distribution of the generated process, this naturally leads to mean-field control problems.

Generative modeling through optimal control

Generative modeling can be viewed as an **optimization problem over probability distributions**: one seeks a mechanism whose output law matches the distribution of the observed data.

- For time series, matching only marginal distributions is generally not enough: the generator should reproduce temporal, path-dependent, and volatility features.
- Stochastic control provides a natural way to **steer the law of a stochastic process** while retaining an interpretable dynamical model.
- When the objective depends explicitly on the distribution of the generated process, this naturally leads to mean-field control problems.

Central Question: Can optimal control of mean-field systems provide a principled framework for designing sample-based time-series algorithms?

Generative models for time series: literature review

- **Score-based diffusion models:** Shen and Kwok '23; Takahashi and Mizuno '24; Wang and Ventre '24; Gao, Zha and Zhou '25; Cao, Chen, Han and Xu '26; Guo, Tang and Xu '26; Cont and Danait '26; ...
- **Schrödinger–Bass bridge models:** Hamdouche, Henry-Labordère and Pham '24; Alouadi, Loeper, Marsala, Mazhar and Pham '26; ...
- **GAN-based models:** Yoon, Jarrett and van der Schaar '19; Vuletić, Prenzel and Cucuringu '24; ...
- **Signature-based models:** Ni, Szpruch, Sabate-Vidales, Xiao, Wiese and Liao '21; Gu, Guo, Jacobs, Kaminsky and Li '24; ...

McKean–Vlasov optimal control with path dependence

- **Full path dependence – probabilistic approach:** Buckdahn, Li, Li and Xing '25
  Stochastic Maximum Principle for controlled mean-field FBSDEs with coefficients depending on the law of the whole paths of the state and control.

- **Full path dependence – dynamic programming approach:** Cosso, Gozzi, Kharroubi, Pham and Rosestolato '23
  Dynamic Programming Principle and Master Bellman equation for path-dependent McKean–Vlasov control problems.

McKean–Vlasov optimal control with path dependence

- **Full path dependence – probabilistic approach:** Buckdahn, Li, Li and Xing '25
  Stochastic Maximum Principle for controlled mean-field FBSDEs with coefficients depending on the law of the whole paths of the state and control.

- **Full path dependence – dynamic programming approach:** Cosso, Gozzi, Kharroubi, Pham and Rosestolato '23
  Dynamic Programming Principle and Master Bellman equation for path-dependent McKean–Vlasov control problems.

In this work, we initiate the study of discrete path dependence in McKean–Vlasov optimal control, where the dependence on the past enters through observations at a finite collection of deterministic times

$$0 < t_1 < \cdots < t_N = T.$$

We will focus on the probabilistic approach and characterize optimality through the Stochastic Maximum Principle and mean-field FBSDEs.

- **Stochastic Maximum Principle for McKean–Vlasov Control with Discrete Path Dependence**
  Joint work with Yadh Hafsi and Huyên Pham. Posted soon on ArXiV.

- **Path dependent mean-field control in discrete time with applications to GenAI**
  Joint work with Amine Zitoun and Huyên Pham. Posted soon on ArXiV.

- **Deep-MKV-TS: Path-Dependent McKean–Vlasov Control for Financial Time Series Generation**
  Joint work with Samer El Boustany, Théo Basseras, Alexandre Alouadi, Yadh Hafsi and Huyên Pham. Preprint available on arXiv:2608.19394.

Outline

Notation

Probability and processes

- $(\Omega, \mathcal{F}, \mathbb{F}, \mathbb{P})$: filtered probability space.
- $(W_t)_{0 \le t \le T}$: Brownian motion.
- $\mathbb{S}^p([0,T];\mathbb{R}^d)$: cadlàg adapted processes such that

$$\mathbb{E}\left[\sup_{0 \le t \le T} |X_t|^p\right] < \infty.$$

- $\mathbb{H}^2([0,T];\mathbb{R}^d)$: progressively measurable processes such that

$$\mathbb{E}\left[\int_0^T |Z_t|^2\, dt\right] < \infty.$$

Laws and path variables

- $\mathcal{P}_2(\mathbb{R}^d)$: probability measures with finite second moment.
- $\mathbb{P}_{X_t}$: law of X_t under $\mathbb{P}$.
- $\Pi := \{0 < t_1 < \ldots < t_N = T\}$: Observation grid.
- $\eta(t)$: number of observation times available at time t.

$$\eta(t) := \sup\{i \in [\![1, N]\!] : t_i \le t\},$$

with convention $\eta(t) = 0$ for $t < t_1$.

- $\boldsymbol{X}_{\eta(t)} = (X_{t_1}, \ldots, X_{t_{\eta(t)}})$: observed path history.
- $\boldsymbol{\mu}_{\eta(t)} = \mathbb{P}_{\boldsymbol{X}_{\eta(t)}}$ observed path law history.

Convention: **Bold symbols** denote path-valued quantities.

Π-non-anticipative maps

Let

$$h : [0, T] \times (\mathbb{R}^d)^N \times \mathcal{P}_2((\mathbb{R}^d)^N) \to \mathbb{R}^{p\times q}.$$

Definition : Π-non-anticipativity

We say that h is Π-*non-anticipative* if there exist measurable maps $h_0, h_1, \dots, h_N$ such that

$$h(t, \boldsymbol{x}, \boldsymbol{\mu}) = \begin{cases} h_0(t), & t \in [t_0, t_1), \\ h_i(t, \boldsymbol{x}_i, \boldsymbol{\mu}_i), & t \in [t_i, t_{i+1}), \quad i = 1, \dots, N-1, \\ h_N(\boldsymbol{x}, \boldsymbol{\mu}), & t = T. \end{cases}$$

Here,

$\boldsymbol{x}_i = (x_1, \dots, x_i)$, $\boldsymbol{\mu}_i = \mathrm{pr}_{1:i}\sharp\boldsymbol{\mu}$ denotes the corresponding marginal law with $\mathrm{pr}_{1:i}(\boldsymbol{x}) := \boldsymbol{x}_i$

Π-non-anticipative maps

Let

$$h : [0, T] \times (\mathbb{R}^d)^N \times \mathcal{P}_2((\mathbb{R}^d)^N) \to \mathbb{R}^{p \times q}.$$

Definition : Π-non-anticipativity

We say that h is Π-*non-anticipative* if there exist measurable maps $h_0, h_1, \ldots, h_N$ such that

$$h(t, \boldsymbol{x}, \boldsymbol{\mu}) = \begin{cases} h_0(t), & t \in [t_0, t_1), \\ h_i(t, \boldsymbol{x}_i, \boldsymbol{\mu}_i), & t \in [t_i, t_{i+1}), \quad i = 1, \ldots, N-1, \\ h_N(\boldsymbol{x}, \boldsymbol{\mu}), & t = T. \end{cases}$$

Here,

$$\boldsymbol{x}_i = (x_1, \ldots, x_i), \qquad \boldsymbol{\mu}_i = \mathrm{pr}_{1:i} \sharp \boldsymbol{\mu} \text{ denotes the corresponding marginal law with } \mathrm{pr}_{1:i}(\boldsymbol{x}) := \boldsymbol{x}_i$$

Interpretation: at time $t \in [t_i, t_{i+1})$, h only depends on the information available up to t_i.

A new class of controlled diffusion processes

- The stochastic dynamics is given by

$$\begin{cases} \mathrm{d}X_t = b(t, \boldsymbol{X}, \mathbb{P}_{\boldsymbol{X}}, X_t, \mathbb{P}_{X_t}, \alpha_t)\mathrm{d}t + \sigma(t, \boldsymbol{X}, \mathbb{P}_{\boldsymbol{X}}, X_t, \mathbb{P}_{X_t}, \alpha_t)\mathrm{d}W_t, \\ X_0 = \xi, \end{cases} \tag{1}$$

  where $t \in [0, T]$, $\boldsymbol{X} = (X_{t_1}, \ldots, X_{t_N})$ and for some Π-non-anticipative maps b, σ satisfying usual assumptions.

- The cost functional is given by

$$J(\alpha) := \mathbb{E}\Big[\int_0^T f(t, \boldsymbol{X}, \mathbb{P}_{\boldsymbol{X}}, X_t, \mathbb{P}_{X_t}, \alpha_t)\mathrm{d}t + g(\boldsymbol{X}, \mathbb{P}_{\boldsymbol{X}})\Big], \tag{2}$$

  where f is a Π-non-anticipative map and g some measurable map.

Well-posedness: Existence and unicity of (1) is guaranteed by standard arguments by solving recursively the SDE on each interval $[t_i, t_{i+1})$ starting from a standard McKean-Vlasov equation on $[t_0, t_1)$.

A new adjoint characterization of the mean-field FBSDE

The Hamiltonian associated with the stochastic control problem (1)–(2) is defined by

$$H(t, \boldsymbol{x}, \boldsymbol{\mu}, x, \mu, y, z, a) := b(t, \boldsymbol{x}, \boldsymbol{\mu}, x, \mu, a) \cdot y + \sigma(t, \boldsymbol{x}, \boldsymbol{\mu}, x, \mu, a) \cdot z + f(t, \boldsymbol{x}, \boldsymbol{\mu}, x, \mu, a).$$

H is a Π-non-anticipative map.

Definition : Adjoint processes

Given ξ an admissible initial condition and α an admissible control, we call

$$(Y, Z) \in \mathbb{S}^2([0, T]; \mathbb{R}^d) \times \mathbb{H}^2([0, T]; \mathbb{R}^{d \times n})$$

adjoint processes of X if

$$Y_t = Y_T + \int_t^T F_s \, \mathrm{d}s - \int_t^T Z_s \, \mathrm{d}W_s + \sum_{k:\, t < t_k \leq t_{N-1}} \xi_k, \qquad t \in [0, T], \quad \mathbb{P} - \text{a.s.} \tag{3}$$

Adjoint coefficients and jump terms

- The continuous driver is

$$F_s := \partial_x H\Big(s, \boldsymbol{X}, \mathbb{P}_{\boldsymbol{X}}, X_s, \mathbb{P}_{X_s}, Y_s, Z_s, \alpha_s\Big) + \widetilde{\mathbb{E}}\Big[\partial_{\tilde{x}} \frac{\delta}{\delta m} H\Big(s, \widetilde{\boldsymbol{X}}, \mathbb{P}_{\boldsymbol{X}}, \widetilde{X}_s, \mathbb{P}_{X_s}, \widetilde{Y}_s, \widetilde{Z}_s, \widetilde{\alpha}_s\Big)(X_s)\Big]. \tag{4}$$

Adjoint coefficients and jump terms

- The continuous driver is

$$F_s := \partial_x H\Big(s, \boldsymbol{X}, \mathbb{P}_{\boldsymbol{X}}, X_s, \mathbb{P}_{X_s}, Y_s, Z_s, \alpha_s\Big) + \widetilde{\mathbb{E}}\Big[\partial_{\tilde{x}} \frac{\delta}{\delta m} H\Big(s, \widetilde{\boldsymbol{X}}, \mathbb{P}_{\boldsymbol{X}}, \widetilde{X}_s, \mathbb{P}_{X_s}, \widetilde{Y}_s, \widetilde{Z}_s, \widetilde{\alpha}_s\Big)(X_s)\Big]. \tag{4}$$

- For $k \in [\![1, N-1]\!]$, the jump terms are

$$\begin{aligned}\xi_k := \mathbb{E}\Bigg[\int_{t_k}^{T} \Bigg(&\partial_{x_k} H\Big(t, \boldsymbol{X}, \mathbb{P}_{\boldsymbol{X}}, X_t, \mathbb{P}_{X_t}, Y_t, Z_t, \alpha_t\Big) + \widetilde{\mathbb{E}}\Big[\partial_{\tilde{x}_k} \frac{\delta}{\delta m} H\Big(t, \widetilde{\boldsymbol{X}}, \mathbb{P}_{\boldsymbol{X}}, \widetilde{X}_t, \mathbb{P}_{X_t}, \widetilde{Y}_t, \widetilde{Z}_t, \widetilde{\alpha}_t\Big)(\boldsymbol{X})\Big]\Bigg) \mathrm{d}t \\ &+ \partial_{x_k} g(\boldsymbol{X}, \mathbb{P}_{\boldsymbol{X}}) + \widetilde{\mathbb{E}}\Big[\partial_{\tilde{x}_k} \frac{\delta}{\delta m} g\Big(\widetilde{\boldsymbol{X}}, \mathbb{P}_{\boldsymbol{X}}\Big)(\boldsymbol{X})\Big] \,\Bigg|\, \mathcal{F}_{t_k}\Bigg].\end{aligned} \tag{5}$$

- The terminal condition is

$$Y_T = \partial_{x_N} g(\boldsymbol{X}, \mathbb{P}_{\boldsymbol{X}}) + \widetilde{\mathbb{E}}\Big[\partial_{\tilde{x}_N} \frac{\delta}{\delta m} g\Big(\widetilde{\boldsymbol{X}}, \mathbb{P}_{\boldsymbol{X}}\Big)(\boldsymbol{X})\Big].$$

Here, $\tilde{\mathbb{E}}$ denotes the expectation over $(\tilde{\Omega}, \tilde{\mathcal{F}}, \tilde{\mathbb{P}})$, on which independent copies of the processes are defined.

Interpretation: Each jump term ξ_k aggregates conditional sensitivity of future costs with respect to the corresponding observed state and its distribution.

A necessary condition for optimality

Under some regularity assumptions on b, σ, f and g and their derivatives, we can state the following necessary condition for optimality.

Theorem : A necessary condition

Suppose that H is a convex map in its control argument for any fixed $(t, \boldsymbol{x}, \boldsymbol{\mu}, x, \mu, y, z)$. Let $(\hat{\alpha}_t)_{0 \leq t \leq T}$ be an optimal control and $(X_t, Y_t, Z_t)_{0 \leq t \leq T}$ be the associated forward and adjoint processes. Then, we have

$$H(t, \boldsymbol{X}, \mathbb{P}_{\boldsymbol{X}}, X_t, \mathbb{P}_{X_t}, Y_t, Z_t, \hat{\alpha}_t) \leq H(t, \boldsymbol{X}, \mathbb{P}_{\boldsymbol{X}}, X_t, \mathbb{P}_{X_t}, Y_t, Z_t, a) \quad \mathrm{d}t \otimes \mathrm{d}\mathbb{P} - \text{a.e},$$

for every $a \in A$.

Definition : Convexity for path-arguments maps

A Π-non anticipative differentiable map $L : [0, T] \times (\mathbb{R}^d)^N \times \mathcal{P}_2((\mathbb{R}^d)^N) \to \mathbb{R}$ is said to be convex if

$$L(t, \boldsymbol{x}', \boldsymbol{\mu}') - L(t, \boldsymbol{x}, \boldsymbol{\mu}) \geq \sum_{i=1}^{\eta(t)} \partial_{x_i} L(t, \boldsymbol{x}, \boldsymbol{\mu}) \cdot (x_i' - x_i) + \mathbb{E}\Big[\sum_{i=1}^{\eta(t)} \partial_{\tilde{x}_i} \frac{\delta}{\delta \mathrm{m}} H(t, \boldsymbol{x}, \boldsymbol{\mu})(\boldsymbol{X}) \cdot (X_i' - X_i) \Big], \tag{6}$$

for any $\boldsymbol{x} = (x_1, \dots, x_N), \boldsymbol{x}' = (x_1', \dots, x_N') \in (\mathbb{R}^d)^N$, $\boldsymbol{\mu}, \boldsymbol{\mu}' \in \mathcal{P}_2((\mathbb{R}^d)^N)$ and for any coupling, i.e., for any pair of random vectors $(\boldsymbol{X}, \boldsymbol{X}') = ((X_1, \dots, X_N), (X_1', \dots, X_N'))$ defined on a common probability space $(\Omega, \mathcal{F}, \mathbb{P})$ with marginal laws $\boldsymbol{X} \sim \boldsymbol{\mu}$, $\boldsymbol{X}' \sim \boldsymbol{\mu}'$.

A sufficient condition for optimality

A necessary condition for optimality can be obtained through the introduced convexity definition as follows.

Theorem : A sufficient condition

Let $(\hat{\alpha}_t)_{0\leq t\leq T} \in \mathcal{A}$ and $(X_t, Y_t, Z_t)_{0\leq t\leq T}$ be the associated controlled state and adjoint processes. If we assume

- $(\mathbb{R}^d)^N \times \mathcal{P}_2((\mathbb{R}^d)^N) \ni (\boldsymbol{x}, \boldsymbol{\mu}) \mapsto g(\boldsymbol{x}, \boldsymbol{\mu})$ is convex in the sense of Eq (6).
- $(\mathbb{R}^d)^N \times \mathcal{P}_2((\mathbb{R}^d)^N) \times \mathbb{R}^d \times \mathcal{P}_2(\mathbb{R}^d) \times A \ni (\boldsymbol{x}, \boldsymbol{\mu}, x, \mu, a) \mapsto H(t, \boldsymbol{x}, \boldsymbol{\mu}, x, \mu, Y_t, Z_t, a)$ is convex $\mathrm{d}t \otimes \mathrm{d}\mathbb{P} -$ a.e in the sense of Eq (6).
- If furthermore

$$H(t, \boldsymbol{X}, \mathbb{P}_{\boldsymbol{X}}, X_t, \mathbb{P}_{X_t}, Y_t, Z_t, \hat{\alpha}_t) = \inf_{a\in A} H(t, \boldsymbol{X}, \mathbb{P}_{\boldsymbol{X}}, X_t, \mathbb{P}_{X_t}, Y_t, Z_t, a), \quad \mathrm{d}t \otimes \mathrm{d}\mathbb{P} - \text{a.e},$$

then $\hat{\alpha}$ is an optimal control.

Solvability of the mean-field FBSDE : Path-dependent LQ systems

- Dynamics of the state process

$$\begin{cases} \mathrm{d}X_t &= \Big[\beta + \sum_{j=1}^{\eta(t)} B_j X_{t_j} + \sum_{j=1}^{\eta(t)} \bar{B}_j \mathbb{E}[X_{t_j}] + CX_t + \bar{C}\mathbb{E}[X_t] + D\alpha_t\Big]\mathrm{d}t + \sigma \mathrm{d}W_t, \\ X_0 &= \xi, \end{cases}$$

- Cost functional

$$\begin{aligned} J(\alpha) =& \mathbb{E}\Big[\int_0^T \Big(P_{\eta(t)} \cdot (\boldsymbol{X}_{\eta(t)}^\top \boldsymbol{X}_{\eta(t)}) + \bar{P}_{\eta(t)} \cdot (\mathbb{E}[\boldsymbol{X}_{\eta(t)}]^\top \mathbb{E}[\boldsymbol{X}_{\eta(t)}]) + \tfrac{1}{2}\alpha_t \cdot R\alpha_t\Big)\mathrm{d}t \\ & + H \cdot (\boldsymbol{X}^\top \boldsymbol{X}) + \bar{H} \cdot (\mathbb{E}[\boldsymbol{X}]^\top \mathbb{E}[\boldsymbol{X}])\Big], \end{aligned}$$

where $\boldsymbol{X}$ is viewed as the $d \times N$ matrix with i-th column X_{t_i} and $\boldsymbol{X}_{\eta(t)}$ as its $d \times \eta(t)$ truncation.

- Ansatz for the adjoint process Y of the form

$$Y_t = \sum_{i=1}^{\eta(t)} \Lambda_i(t) X_{t_i} + \sum_{i=1}^{\eta(t)} \bar{\Lambda}_i(t)\mathbb{E}[X_{t_i}] + \Gamma(t)X_t + \bar{\Gamma}(t)\mathbb{E}[X_t] + \chi(t), \quad 0 \le t \le T, \tag{7}$$

where $(\Lambda_i, \bar{\Lambda}_i) : [t_i, T] \to \mathbb{R}^{d\times d}$, $(\Gamma, \bar{\Gamma}) : [0, T] \to \mathbb{R}^{d\times d}$ and $\chi : [0, T] \to \mathbb{R}^d$ are deterministic maps for any $1 \le i \le N$ with convention that $\Gamma(T) = \bar{\Gamma}(T) = 0$, obtained through Riccati equations and jumps.

Solvability of the mean-field FBSDE : Path-dependent LQ systems

- Dynamics of the state process

$$\begin{cases} \mathrm{d}X_t &= \Big[\beta + \sum_{j=1}^{\eta(t)} B_j X_{t_j} + \sum_{j=1}^{\eta(t)} \bar{B}_j \mathbb{E}[X_{t_j}] + CX_t + \bar{C}\mathbb{E}[X_t] + D\alpha_t\Big]\mathrm{d}t + \sigma \mathrm{d}W_t, \\ X_0 &= \xi, \end{cases}$$

- Cost functional

$$\begin{aligned} J(\alpha) =& \mathbb{E}\Big[\int_0^T \Big(P_{\eta(t)} \cdot (\boldsymbol{X}_{\eta(t)}^\top \boldsymbol{X}_{\eta(t)}) + \bar{P}_{\eta(t)} \cdot (\mathbb{E}[\boldsymbol{X}_{\eta(t)}]^\top \mathbb{E}[\boldsymbol{X}_{\eta(t)}]) + \tfrac{1}{2}\alpha_t \cdot R\alpha_t\Big)\mathrm{d}t \\ & + H \cdot (\boldsymbol{X}^\top \boldsymbol{X}) + \bar{H} \cdot (\mathbb{E}[\boldsymbol{X}]^\top \mathbb{E}[\boldsymbol{X}])\Big], \end{aligned}$$

  where $\boldsymbol{X}$ is viewed as the $d \times N$ matrix with i-th column X_{t_i} and $\boldsymbol{X}_{\eta(t)}$ as its $d \times \eta(t)$ truncation.
- Ansatz for the adjoint process Y of the form

$$Y_t = \sum_{i=1}^{\eta(t)} \Lambda_i(t) X_{t_i} + \sum_{i=1}^{\eta(t)} \bar{\Lambda}_i(t)\mathbb{E}[X_{t_i}] + \Gamma(t)X_t + \bar{\Gamma}(t)\mathbb{E}[X_t] + \chi(t), \quad 0 \le t \le T, \tag{7}$$

  where $(\Lambda_i, \bar{\Lambda}_i) : [t_i, T] \to \mathbb{R}^{d\times d}$, $(\Gamma, \bar{\Gamma}) : [0, T] \to \mathbb{R}^{d\times d}$ and $\chi : [0, T] \to \mathbb{R}^d$ are deterministic maps for any $1 \le i \le N$ with convention that $\Gamma(T) = \bar{\Gamma}(T) = 0$, obtained through Riccati equations and jumps.

Takeaway: Existence and uniqueness can be shown for path-dependent linear quadratic systems relying on (7) and the stability of the process $\big(\mathbb{E}[X_t|\mathcal{F}_{t_k}]\big)_{t\ge t_k}$ for any $1 \le k \le N-1$.

Solvability of the mean-field FBSDE : Time series generation

- Dynamics of the state process

$$\begin{cases} \mathrm{d}X_t & = \alpha_t \mathrm{d}t + \sigma(t, X_t)\mathrm{d}W_t, \\ X_0 & = \xi \end{cases}$$

- Given a penalty parameter $\lambda > 0$, we define the cost functional as

$$J^\lambda(\alpha) := \mathbb{E}\Big[\int_0^T \frac{1}{2}|\alpha_t|^2 \mathrm{d}t\Big] + \lambda \mathrm{MMD}_K^2(\mathbb{P}_{(X_{t_1},\ldots,X_{t_N})}|\rho),$$

where the Maximum Mean Discrepancy (MMD) map is defined as $\mathrm{MMD}_K(\cdot|\rho)$

$$\mathrm{MMD}_K^2(\mu|\rho) = \frac{1}{2}\int_{((\mathbb{R}^d)^N)^2} K(\boldsymbol{x}, \boldsymbol{y})(\mu - \rho)(\mathrm{d}\boldsymbol{x})(\mu - \rho)(\mathrm{d}\boldsymbol{y}),$$

for some differentiable kernel map $K : (\mathbb{R}^d)^N \times (\mathbb{R}^d)^N \to \mathbb{R}$ for which $D_K(\mu|\rho) = 0$ i.f.f. $\mu = \rho$.

Takeaway: Working with MMD here is particularly convenient as the whole regularity requirements for the contraction argument is obtained through the regularity of the kernel K itself.

Solvability of the mean-field FBSDE : Time series generation

Following our approach, the MF-FBSDE system $(X^\lambda, Y^\lambda, Z^\lambda)$ is given by

$$\begin{cases} \mathrm{d}X_t^\lambda &= -Y_t^\lambda \mathrm{d}t + \sigma(t, X_t^\lambda)\mathrm{d}W_t, \\ X_0^\lambda &= 0, \\ \mathrm{d}Y_t^\lambda &= -\partial_x\big(\mathrm{Tr}(\sigma(t,X_t^\lambda)^\top Z_t^\lambda)\big)\mathrm{d}t + Z_t^\lambda \mathrm{d}W_t, \quad t_{k-1} \le t < t_k, \quad k \in [\![1, N-1]\!], \quad \text{and } t_{N-1} \le t \le t_N, \\ Y_{t_k}^\lambda - Y_{t_k^-}^\lambda &= -\lambda \mathbb{E}\Big[\int_{(\mathbb{R}^d)^N} \partial_{x_k} K(\boldsymbol{X}^\lambda, \boldsymbol{y})\mathbb{P}_{(X_{t_1}^\lambda,\dots,X_{t_N}^\lambda)}(\mathrm{d}\boldsymbol{y}) - \int_{(\mathbb{R}^d)^N} \partial_{x_k} K(\boldsymbol{X}^\lambda, \boldsymbol{y})\rho(\mathrm{d}\boldsymbol{y}) \Big| \mathcal{F}_{t_k}\Big], \quad k \in [\![1, N-1]\!], \\ Y_{t_N}^\lambda &= \lambda\Big(\int_{(\mathbb{R}^d)^N} \nabla_{x_N} K(\boldsymbol{X}^\lambda, \boldsymbol{y})\mathbb{P}_{(X_{t_1}^\lambda,\dots,X_{t_N}^\lambda)}(\mathrm{d}\boldsymbol{y}) - \int_{(\mathbb{R}^d)^N} \nabla_{x_N} K(\boldsymbol{X}^\lambda, \boldsymbol{y})\rho(\mathrm{d}\boldsymbol{y})\Big). \end{cases} \tag{8}$$

Interpretation: Since J^λ is finite for any $\lambda > 0$, one expects $\mathrm{MMD}_K(\mathbb{P}_{(X_{t_1}^\lambda,\dots,X_{t_N}^\lambda)}|\rho) \underset{\lambda\to\infty}{\to} 0$, where $(X^\lambda, Y^\lambda, Z^\lambda)$ is a solution to (8).

Theorem : Existence and uniqueness in small time

Let $K \in \mathcal{C}^2\big((\mathbb{R}^d)^N \times (\mathbb{R}^d)^N\big)$ with uniformly bounded derivative maps $\partial_x K$, $\partial_{xx} K$ and $\partial_y \partial_x K$. Suppose that σ has a uniformly bounded and Lipschitz derivative $\partial_x \sigma$, and that there exists a map $\Sigma := (\Sigma_1, \dots, \Sigma_n) : [0, T] \times \mathbb{R}^d \to \mathbb{R}^n$ such that

$$\partial_x \sigma^j(t, x) = \Sigma_j(t, x) I_d, \qquad j = 1, \dots, n,$$

where σ^j denotes the j-th column of σ and I_d is the identity matrix of $\mathbb{R}^{d \times d}$.
Then, there exists $\delta > 0$ depending on the bounds of $\partial_x \sigma$, the Lipschitz constants of $\nabla_{\boldsymbol{x}} K$, λ, and the number of points N in the grid, such that for any $T \in (0, \delta]$, the MF-FBSDE (8) admits a unique solution

$$(X^\lambda, Y^\lambda, Z^\lambda) \in \mathbb{S}^4([0, T]; \mathbb{R}^d) \times \mathbb{S}^\infty([0, T]; \mathbb{R}^d) \times \mathbb{H}^2([0, T]; \mathbb{R}^{d \times n}).$$

Outline

Introduction and motivation

Path-dependent Mc-Kean Vlasov optimal control

Financial time series generation

Synthetic experiments : Choice of the generative model

Question: Can our mean-field optimal control approach adjust a reference model by matching selected path and volatility features of generated scenarios to those observed in the data ?

Synthetic experiments : Choice of the generative model

Question: Can our mean-field optimal control approach adjust a reference model by matching selected path and volatility features of generated scenarios to those observed in the data ?

Starting from some reference model, fitted from the data, from which one can compute drift and volatility Π-non anticipative maps b^{ref} and σ^{ref}, we consider the following optimization problem

- State dynamics

$$\mathrm{d}X_t = b^{\text{ref}}(t, \boldsymbol{X})\mathrm{d}t + \sigma_t \mathrm{d}W_t, \tag{9}$$

- Cost functional

$$\begin{aligned} J(\sigma) = \mathbb{E}\Big[\int_0^T \mathcal{E}\big(\sigma_t, \sigma^{\text{ref}}(t, \boldsymbol{X})\big)\mathrm{d}t\Big] &+ \frac{w_{\text{path}}}{2}\mathrm{MMD}_K^2(\mathbb{P}_{(X_{t_1},\dots,X_{t_N})}|\rho) \\ &+ \frac{w_{\text{vol}}}{2}\mathrm{MMD}_K^2(\varphi_{\text{vol}}\sharp\mathbb{P}_{(X_{t_1},\dots,X_{t_N})}|\varphi_{\text{vol}}\sharp\rho), \end{aligned} \tag{10}$$

where for any $\sigma, \bar{\sigma} \in \mathbb{S}^d_{++}$, $\mathcal{E}(\sigma, \bar{\sigma}) := \mathrm{Tr}\big(\bar{\sigma}^{-1}(\sigma - \bar{\sigma})\big) - \log\Big(\frac{\det(\sigma)}{\det(\bar{\sigma})}\Big)$, is the specific entropy, φ_{vol} is a map specifying the targeted stylized facts that we aim to reproduce and $w_{\text{path}}, w_{\text{vol}} > 0$ balance the relative importance of these two objectives.

Optimal control form: Following the approach developed in this talk, a good optimal control candidate is given by $\sigma_t^\star = \Big(\sigma^{\text{ref}}(t, \boldsymbol{X}^\star) + Z_t^\star)\Big)^{-1}$ for any $0 \le t \le T$.

Numerical implementation: learning the adjoint process

For the numerical experiments, we use the discrete-time formulation of the control problem (9)-(10)

- **Time discretization.** We discretize $[0, T]$ over the grid

$$\Pi_{\mathrm{disc}} = \{0 = s_0 < s_1 < \cdots < s_M = T\}.$$

- **Parametrization of the adjoint process.** The optimal martingale component $(Z^{\star}_{s_k})_{k=0}^{M-1}$ is approximated by a recurrent neural network:

$$Z^{\star}_{s_k} \approx \widehat{Z}^{\theta}_{s_k} := \mathrm{GRU}_{\theta}(\mathcal{H}_{s_k}),$$

where $\mathcal{H}_{s_k}$ denotes the information available up to time s_k.

- **Adjoint learning.** At each training iteration, proxy targets $Z^{\mathrm{proxy}}_{s_k}$ are constructed from the discrete-time adjoint equations, and the network is trained through

$$\mathcal{L}_{\mathrm{adj}}(\theta) = \frac{1}{PM} \sum_{p=1}^{P} \sum_{k=0}^{M-1} \left\| \widehat{Z}^{\theta,p}_{s_k} - Z^{\mathrm{proxy},p}_{s_k} \right\|^2 .$$

- **Optimization.** The network parameters are updated using AdamW, and after the final training iteration K,

$$Z^{\star}_{s_k} \approx \widehat{Z}^{\theta^{(K)}}_{s_k}.$$

Recovering hidden volatility from a mixture of Heston models

- **Setup.** Target: mixture of 8 Heston regimes with different (θ, ξ, ρ). Baseline: a single Heston model fitted to the pooled data.
- **Distributional evaluation.** All diagnostics compare laws of path-dependent features.

$$\phi_{\mathrm{RV}}(x) := \frac{1}{N-1}\sum_{i=1}^{N-1}\frac{r_i^2}{\Delta t_i}, \qquad \mathrm{RV}\ W_1 := \frac{W_1(\phi_{\mathrm{RV}}\sharp\mu_{\mathrm{data}}, \phi_{\mathrm{RV}}\sharp\mu_{\mathrm{gen}})}{\mathbb{E}_{\mu_{\mathrm{data}}}[\phi_{\mathrm{RV}}(X)]}.$$

$$\mathrm{Regime\ TVD} := d_{\mathrm{TVD}}\left(c^\star\sharp\mu_{\mathrm{data}}, c^\star\sharp\mu_{\mathrm{gen}}\right), \qquad c^\star : \text{Class}.$$

Method	SWD	RV W_1	$\lvert r\rvert$ ACF	TVD	MDD W_1
Baseline	0.347	0.479	0.322	0.629	0.309
Deep-MKV-TS	**0.237**	**0.095**	**0.046**	**0.331**	**0.009**
SBTS	0.407	0.493	0.094	0.563	0.478
SBBTS	0.573	0.412	0.207	0.360	0.331

- **RV W_1:** matching of volatility levels.
- **Regime TVD:** recovery of latent volatility regimes.

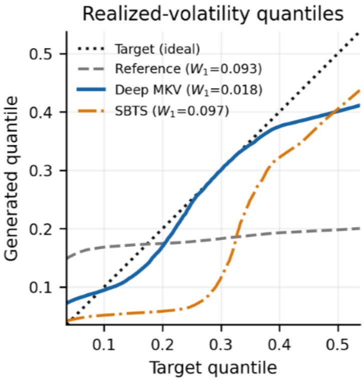

Takeaway: Deep-MKV-TS better captures both the **distribution of volatility levels** and the **latent regime structure**.

From path-dependent McKean-Vlasov optimal control to Generative AI

What we presented

- We introduced a framework based on optimal control for the design of time-series generation algorithms.
- The method incorporates path-dependent and distributional information directly into the modelling objective.
- The resulting control problem can be solved using a stochastic maximum principle and sample-based numerical methods.

Potential improvements

Theoretical side

- Extend the current short-time well-posedness result to arbitrary time horizons.
- Establish existence and uniqueness when the control acts on the volatility, which is particularly relevant for financial applications.
- Derive sufficient optimality conditions: the required convexity assumptions are not available in the current volatility-control setting.
- Go beyond MMD-type objectives and allow for more general distributional costs g.

Potential improvements

Theoretical side

- Extend the current short-time well-posedness result to arbitrary time horizons.
- Establish existence and uniqueness when the control acts on the volatility, which is particularly relevant for financial applications.
- Derive sufficient optimality conditions: the required convexity assumptions are not available in the current volatility-control setting.
- Go beyond MMD-type objectives and allow for more general distributional costs g.

Numerical side

- Improve the stability of the forward–backward learning procedure.
- Establish convergence guarantees for the algorithm in relevant tractable model classes.
- Improve scalability for high-dimensional and long-horizon time series.

Perspective: The framework remains flexible: rather than matching the full distribution directly, the cost functional can target problem-specific path and volatility features.

Appendix

Alternative formulation of the mean-field FBSDE

The BSDE (3) can be alternatively formulated as a coupled system of N continuous backward equations. Namely, we define over the interval $[t_{N-1}, t_N]$ the BSDE

- $$\begin{cases} \mathrm{d}Y_t^{(N)} & = \int_{t_{N-1}}^t F_s^{(N)} \mathrm{d}s - \int_t^T Z_s^{(N)} \mathrm{d}W_s, \quad t_{N-1} \le t \le t_N \\ Y_T^{(N)} & = \partial_{x_N} g(\boldsymbol{X}, \mathbb{P}_{\boldsymbol{X}}) + \tilde{\mathbb{E}}\left[\partial_{\tilde{x}_N} \frac{\delta}{\delta \mathbf{m}} g(\tilde{\boldsymbol{X}}, \mathbb{P}_{\boldsymbol{X}})(\boldsymbol{X})\right], \end{cases}$$

- we then proceed recursively by backward induction in time by solving $(Y^{(k)}, Z^{(k)})$ on $[t_{k-1}, t_k]$ for $k \in [\![1, N-1]\!]$

  $$\begin{cases} \mathrm{d}Y_t^{(k)} & = \int_{t_{k-1}}^t F_s^{(k)} \mathrm{d}s - \int_{t_{k-1}}^t Z_s^{(k)} \mathrm{d}W_s, \quad t_{k-1} \le t \le t_k, \\ Y_{t_k}^{(k)} & = Y_{t_k}^{(k+1)} + \xi_k^{(k+1)} \end{cases}$$

  where the process $(F_s^{(k)})_{t_{k-1} \le s \le t_k}$ is obtained by replacing (Y_s, Z_s) in (4) by $(Y_s^{(k)}, Z_s^{(k)})$ and where $\xi_k^{(k+1)}$ is given by (5) where we replaced (Y_t, Z_t) by $(Y_t^{(l+1)}, Z_t^{(l+1)})$ on each subinterval $t_l \le t \le t_{l+1}$ for $k \le l \le N-1$.

- Global backward equation is given by

$$\begin{cases} Y_t & = \sum_{k=1}^N \mathbb{1}_{t_{k-1} \le t < t_k} Y_t^{(k)} + \mathbb{1}_{t = t_N} Y_T^{(N)}, \quad 0 \le t \le T, \quad \mathbb{P}-\text{a.s} \\ Z_t & = \sum_{k=1}^N \mathbb{1}_{t_{k-1} \le t < t_k} Z_t^{(k)}, \quad \mathrm{d}t \otimes \mathrm{d}\mathbb{P} - \text{a.e.} \end{cases} \tag{11}$$

Extension to conditional generation

- **Conditional setting.** Let $I = (I_t)_{0 \le t \le T}$ be an exogenous finite-state process, e.g. encoding observable financial regimes, and
$$\mathcal{F}_t := \sigma\big(\xi, (W_s, I_s)_{0 \le s \le t}\big)\,.$$
The objective is to learn the conditional distribution
$$\omega \longmapsto \mathbb{P}_{\boldsymbol{X} \mid I = I(\omega)}, \qquad \boldsymbol{X} = (X_{t_1}, \ldots, X_{t_N}).$$

- **Control problem.**
$$\mathrm{d}X_t = \alpha_t \,\mathrm{d}t + \sigma(t, X_t, I_t)\,\mathrm{d}W_t, \qquad X_0 = \xi,$$
$$J(\alpha) := \mathbb{E}\left[\int_0^T \frac{1}{2}|\alpha_t|^2 \,\mathrm{d}t + \lambda\, g\Big(\mathbb{P}^I_{\boldsymbol{X}} \mid \rho^I\Big)\right], \qquad \mathbb{P}^I_{\boldsymbol{X}} := \mathbb{P}(\boldsymbol{X} \in \cdot \mid I).$$

- **Associated MF-FBSDE.** If I is a finite-state Markov chain, the optimality system formally reads
$$\begin{cases} \mathrm{d}X_t^\lambda = -Y_t^\lambda \,\mathrm{d}t + \sigma(t, X_t^\lambda, I_t)\,\mathrm{d}W_t, \\ \mathrm{d}Y_t^\lambda = -\partial_x\big(\mathrm{Tr}(\sigma(t, X_t^\lambda, I_t)^\top Z_t^\lambda)\big), \mathrm{d}t + Z_t^\lambda \,\mathrm{d}W_t + \sum\limits_{j \ne I_{t-}} U_t^{I_{t-},j}\,\mathrm{d}M_t^{I_{t-},j}, \\ \Delta Y_{t_k}^\lambda = -\lambda\,\mathbb{E}\left[\partial_{\tilde{x}_k}\frac{\delta g}{\delta m}(\mathbb{P}^I_{\boldsymbol{X}^\lambda} \mid \rho^I)(\boldsymbol{X}^\lambda)\Big|\,\mathcal{F}_{t_k}\right], \quad k = 1, \ldots, N-1, \\ Y_T^\lambda = \lambda\,\partial_{\tilde{x}_N}\frac{\delta g}{\delta m}(\mathbb{P}^I_{\boldsymbol{X}^\lambda} \mid \rho^I)(\boldsymbol{X}^\lambda). \end{cases}$$

Takeaway: Conditioning on I yields a **regime-adaptive generator** while preserving the path-dependent mean-field control structure.

Discrete-time framework

The continuous-time framework can be extended to a genuinely discrete-time setting, without restricting the dynamics to an Euler discretization of an Itô process.

- **Time grids.** We distinguish the evolution grid

  $$\Pi_{\mathrm{disc}} := \{0 = s_0 < s_1 < \cdots < s_M = T\},$$

  and the observation times

  $$\Pi_{\mathrm{obs}} := \{0 < t_1 < \cdots < t_N = T\} \subset \Pi_{\mathrm{disc}}.$$

  We set

  $$\eta(s_k) := \max\{j : t_j \le s_k\}, \qquad \boldsymbol{X}_{\eta(s_k)} := (X_{t_1}, \dots, X_{t_{\eta(s_k)}}).$$

- **State dynamics.**

  $$\begin{cases} X_{s_{k+1}} = F_k\Big(X_{s_k}, \mathbb{P}_{X_{s_k}}, \boldsymbol{X}_{\eta(s_k)}, \mathbb{P}_{\boldsymbol{X}_{\eta(s_k)}}, \alpha_{s_k}, \epsilon_{s_{k+1}}\Big), \\ X_{s_0} = \xi, \end{cases} \qquad k = 0, \dots, M-1.$$

- **Cost functional.**

  $$J(\alpha) := \mathbb{E}\left[\sum_{k=0}^{M-1} f_k\Big(X_{s_k}, \mathbb{P}_{X_{s_k}}, \boldsymbol{X}_{\eta(s_k)}, \mathbb{P}_{\boldsymbol{X}_{\eta(s_k)}}, \alpha_{s_k}\Big) + g(\boldsymbol{X}, \mathbb{P}_{\boldsymbol{X}})\right].$$

Key point: The states X_{s_k} describe the full discrete-time evolution, while $X_{t_1}, \dots, X_{t_N}$ are the observations entering the path-dependent objective.

Discrete-time Hamiltonian

- Given the adjoint variable $Y_{s_{k+\mathbf{1}}}$, we introduce the family of random Hamiltonian maps

$$
\begin{aligned}
H_k(\omega, x, \mu, \boldsymbol{x}, \boldsymbol{\mu}, a) := {} & f_k(x, \mu, \boldsymbol{x}, \boldsymbol{\mu}, a) \\
& + \mathbb{E}_{s_k}\Big[Y_{s_{k+\mathbf{1}}} \cdot F_k(x, \mu, \boldsymbol{x}, \boldsymbol{\mu}, a, \epsilon_{s_{k+\mathbf{1}}}) \Big](\omega).
\end{aligned}
$$

- In particular, along the controlled trajectory,

$$
H_k := H_k\Big(X_{s_k}, \mathbb{P}_{X_{s_k}}, \boldsymbol{X}_{\eta(s_k)}, \mathbb{P}_{\boldsymbol{X}_{\eta(s_k)}}, \alpha_{s_k} \Big).
$$

- The Hamiltonian therefore incorporates both the **immediate running cost** and the **future sensitivity propagated by the adjoint process**.

Interpretation: Unlike an Euler discretization of a continuous-time Hamiltonian, H_k is defined directly from the general transition map F_k.

Discrete-time adjoint system and optimal control

For $k = 0, \ldots, M$, the adjoint process is characterized backward by

$$
Y^\star_{s_k} = \begin{cases}
\nabla_{x_N} g\Big(\boldsymbol{X}^\star, \mathbb{P}_{\boldsymbol{X}^\star}\Big) + \widetilde{\mathbb{E}}\Big[\nabla_{\tilde{x}_N} \dfrac{\delta g}{\delta m}\Big(\widetilde{\boldsymbol{X}}^\star, \mathbb{P}_{\boldsymbol{X}^\star}\Big)(\boldsymbol{X}^\star)\Big], & s_k = t_N = T, \\
\partial_x H_k\Big(X^\star_{s_k}, \mathbb{P}_{X^\star_{s_k}}, \boldsymbol{X}^\star_{\eta(s_k)}, \mathbb{P}_{\boldsymbol{X}^\star_{\eta(s_k)}}, \alpha^\star_{s_k}\Big) \\
\qquad + \widetilde{\mathbb{E}}\Big[\nabla_{\tilde{x}} \dfrac{\delta H_k}{\delta m}\Big(\widetilde{X}^\star_{s_k}, \mathbb{P}_{X^\star_{s_k}}, \widetilde{\boldsymbol{X}}^\star_{\eta(s_k)}, \mathbb{P}_{\boldsymbol{X}^\star_{\eta(s_k)}}, \widetilde{\alpha}^\star_{s_k}\Big)(X^\star_{s_k})\Big], & s_k \notin \Pi_{\mathrm{obs}}, \\
\partial_x H_k\Big(X^\star_{t_j}, \mathbb{P}_{X^\star_{t_j}}, \boldsymbol{X}^\star_j, \mathbb{P}_{\boldsymbol{X}^\star_j}, \alpha^\star_{t_j}\Big) + \widetilde{\mathbb{E}}\Big[\nabla_{\tilde{x}} \dfrac{\delta H_k}{\delta m}\Big(\widetilde{X}^\star_{t_j}, \mathbb{P}_{X^\star_{t_j}}, \widetilde{\boldsymbol{X}}^\star_j, \mathbb{P}_{\boldsymbol{X}^\star_j}, \widetilde{\alpha}^\star_{t_j}\Big)(X^\star_{t_j})\Big] \\
\qquad + \mathbb{E}_{t_j}\Big[\displaystyle\sum_{k' : j \le \eta(s_{k'})} \Big(\nabla_{x_j} H_{k'}\Big(X^\star_{s_{k'}}, \mathbb{P}_{X^\star_{s_{k'}}}, \boldsymbol{X}^\star_{\eta(s_{k'})}, \mathbb{P}_{\boldsymbol{X}^\star_{\eta(s_{k'})}}, \alpha^\star_{s_{k'}}\Big) \\
\qquad\qquad + \widetilde{\mathbb{E}}\Big[\nabla_{\tilde{x}_j} \dfrac{\delta H_{k'}}{\delta m}\Big(\widetilde{X}^\star_{s_{k'}}, \mathbb{P}_{X^\star_{s_{k'}}}, \widetilde{\boldsymbol{X}}^\star_{\eta(s_{k'})}, \mathbb{P}_{\boldsymbol{X}^\star_{\eta(s_{k'})}}, \widetilde{\alpha}^\star_{s_{k'}}\Big)(\boldsymbol{X}^\star_{\eta(s_{k'})})\Big]\Big)\Big] \\
\qquad + \mathbb{E}_{t_j}\Big[\nabla_{x_j} g\Big(\boldsymbol{X}^\star, \mathbb{P}_{\boldsymbol{X}^\star}\Big) + \widetilde{\mathbb{E}}\Big[\nabla_{\tilde{x}_j} \dfrac{\delta g}{\delta m}\Big(\widetilde{\boldsymbol{X}}^\star, \mathbb{P}_{\boldsymbol{X}^\star}\Big)(\boldsymbol{X}^\star)\Big]\Big], & s_k = t_j, \quad j < N.
\end{cases}
$$

The optimal control satisfies, for $k = 0, \ldots, M-1$,

$$
\boxed{\alpha^\star_{s_k} \in \operatorname*{argmin}_{a \in A} H_k\Big(X^\star_{s_k}, \mathbb{P}_{X^\star_{s_k}}, \boldsymbol{X}^\star_{\eta(s_k)}, \mathbb{P}_{\boldsymbol{X}^\star_{\eta(s_k)}}, a\Big)}
$$

Definition of the map φ_{vol} (1/2)

Let $\boldsymbol{x} = (x_{t_1}, \ldots, x_{t_N}) \in (\mathbb{R}^d)^N$ be a normalized log-price path, with

$$r_i = x_{t_{i+1}} - x_{t_i}, \qquad i \in [\![1, N-1]\!].$$

Broad path features.

$$\begin{aligned}\phi_{\mathrm{path}}(\boldsymbol{x}) &= (x_{t_1}, \ldots, x_{t_N}),\\ \phi_{\mathrm{terminal}}(\boldsymbol{x}) &= x_{t_N},\\ \phi_{\mathrm{return}}(\boldsymbol{x}) &= (r_i)_{i=1}^{N-1}.\end{aligned}$$

Absolute returns.
We use the differentiable approximation

$$a_i = \sqrt{r_i^2 + \varepsilon_a^2} - \varepsilon_a,$$

which reduces to $|r_i|$ when $\varepsilon_a = 0$. Then

$$\phi_{\mathrm{abs}}(\boldsymbol{x}) = (a_i)_{i=1}^{N-1}.$$

Squared returns.

$$\phi_{\mathrm{sq}}(\boldsymbol{x}) = (r_i^2)_{i=1}^{N-1}.$$

Instantaneous realized variance.
Setting

$$q_i = \frac{r_i^2}{\Delta t_i},$$

we define

$$\phi_{\mathrm{instRV}}(\boldsymbol{x}) = (q_i)_{i=1}^{N-1}.$$

Global realized variance.

$$\phi_{\mathrm{globalRV}}(\boldsymbol{x}) = \frac{1}{N-1} \sum_{i=1}^{N-1} q_i.$$

Definition of the map φ_{vol} (2/2)

Rolling realized volatility.
For $w \in [\![1, N-1]\!]$,

$$v_i^{(w)} = \left(\frac{1}{w} \sum_{j=i-w+1}^{i} r_j^2 + \varepsilon_v^2 \right)^{1/2},$$

and

$$\phi_{\mathrm{rollRV}}^{(w)}(\boldsymbol{x}) = (v_i^{(w)})_{i=w}^{N-1}.$$

State–variance pairs.
To capture how volatility varies with the current state,

$$\phi_{\mathrm{stateRV}}(\boldsymbol{x}) = \big((x_{t_i}, q_i)\big)_{i=1}^{N-1}, \qquad q_i = \frac{r_i^2}{\Delta t_i}.$$

Volatility-persistence features.
For $u_i \in \{a_i, r_i^2\}$, define

$$\widetilde{u}_i = \frac{u_i - \overline{u}}{s_u},$$

where $\overline{u}$ and s_u are the pooled mean and standard deviation.
For $\ell \in \mathcal{L} \subseteq [\![0, N-2]\!]$,

$$\phi_{\mathrm{lag},\ell}^{u}(\boldsymbol{x}) = \big(\widetilde{u}_i \widetilde{u}_{i+\ell}\big)_{i=1}^{N-1-\ell}.$$

Complete volatility map.

$$\begin{aligned}\varphi_{\mathrm{vol}}(\boldsymbol{x}) = \Big(&\phi_{\mathrm{abs}}, \phi_{\mathrm{sq}}, \phi_{\mathrm{instRV}}, \phi_{\mathrm{globalRV}},\\ &\phi_{\mathrm{rollRV}}^{(w)}, \phi_{\mathrm{stateRV}},\\ &(\phi_{\mathrm{lag},\ell}^{a})_{\ell\in\mathcal{L}},\\ &(\phi_{\mathrm{lag},\ell}^{r^2})_{\ell\in\mathcal{L}}\Big)(\boldsymbol{x}).\end{aligned}$$

Definition of evaluation metrics (1/2)

Let μ_{data} and μ_{gen} denote respectively the distributions of real and generated paths.

Sliced Wasserstein distance (SWD).
For $\theta \in \mathbb{S}^{dN-1}$, let $P^{\theta}(x) = \theta \cdot x$. We compute

$$\mathrm{SWD} := \int_{\mathbb{S}^{dN-1}} \mathcal{W}_1 \left(P^{\theta} \sharp \mu_{\text{data}}, P^{\theta} \sharp \mu_{\text{gen}} \right) \, \mathrm{d}\sigma(\theta).$$

It measures the discrepancy between the **full distributions of paths** through one-dimensional projections.

Realized volatility (RV W_1).
For

$$\phi_{\text{globalRV}}(\boldsymbol{x}) = \frac{1}{N-1} \sum_{i=1}^{N-1} \frac{r_i^2}{\Delta t_i},$$

we define

$$\mathrm{RV}\ W_1 := \frac{\mathcal{W}_1 \left(\phi_{\text{globalRV}} \sharp \mu_{\text{data}}, \phi_{\text{globalRV}} \sharp \mu_{\text{gen}} \right)}{\mathbb{E}_{\mu_{\text{data}}}[\phi_{\text{globalRV}}]}.$$

Absolute-return ACF ($|r|$ ACF).
For a lag ℓ, define

$$\rho^{\mu}_{|r|}(\ell) = \frac{1}{N-1-\ell} \sum_{i=1}^{N-1-\ell} \mathrm{Corr}_{\mu} \left(|r_i|, |r_{i+\ell}| \right).$$

For a chosen set of lags $\mathcal{L}$,

$$|r|\ \mathrm{ACF} := \left(\frac{1}{|\mathcal{L}|} \sum_{\ell \in \mathcal{L}} \left(\rho^{\mu_{\text{data}}}_{|r|}(\ell) - \rho^{\mu_{\text{gen}}}_{|r|}(\ell) \right)^2 \right)^{1/2}.$$

It measures whether generated paths reproduce **volatility clustering and persistence**.

Interpretation.

- SWD: full path distribution.
- RV W_1: volatility distribution.
- $|r|$ ACF: volatility persistence.

Definition of evaluation metrics (2/2)

Maximum drawdown (MDD W_1).
For a log-price path $\boldsymbol{x}$, define

$$\mathrm{MDD}(\boldsymbol{x}) = \max_{0 \le n \le N} \left(\max_{0 \le k \le n} x_{t_k} - x_{t_n} \right).$$

We compare its distributions through

$$\mathrm{MDD}\ W_1 := \frac{\mathcal{W}_1\left(\mathrm{MDD}\sharp\mu_{\mathrm{data}}, \mathrm{MDD}\sharp\mu_{\mathrm{gen}}\right)}{\mathbb{E}_{\mu_{\mathrm{data}}}[\mathrm{MDD}]}.$$

This metric measures whether the model reproduces the distribution of **extreme peak-to-trough losses**.

Regime total variation distance (TVD).
For the mixture experiment, let

$$c^\star : (\mathbb{R}^d)^N \longrightarrow [\![1, L]\!]$$

be a trained classifier assigning each path to a mixture regime.
The induced regime distributions are

$$c^\star\sharp\mu_{\mathrm{data}}, \qquad c^\star\sharp\mu_{\mathrm{gen}}.$$

We define

$$\begin{aligned} \mathrm{TVD} &:= d_{\mathrm{TVD}}\left(c^\star\sharp\mu_{\mathrm{data}}, c^\star\sharp\mu_{\mathrm{gen}}\right) \\ &= \frac{1}{2}\sum_{k=1}^{L} \left|\mathbb{P}_{\mu_{\mathrm{data}}}(c^\star = k) - \mathbb{P}_{\mu_{\mathrm{gen}}}(c^\star = k)\right|. \end{aligned}$$

It measures how well the generated data preserve the **mixture/regime proportions**.

Construction of reference maps b^{ref} and σ^{ref} (1/2)

Let

$$X^p_{t_i} = \log\left(\frac{S^p_{t_i}}{S^p_{t_1}}\right), \qquad \Delta X^p_i = X^p_{t_{i+1}} - X^p_{t_i}, \qquad \Delta t_i = t_{i+1} - t_i.$$

Reference drift b^{ref}. We encode the past trajectory through short- and long-memory trend signals

$$\psi^{\mathrm{ref}}_i(\boldsymbol{x}_{1:i}) = \left(T^{(1)}_i, T^{(2)}_i\right) \in \mathbb{R}^{2d}, \qquad q^{\mathrm{ref}}_i = (1, \psi^{\mathrm{ref}}_i),$$

where each component is standardized on the training set. The trend memories satisfy

$$T^{(a)}_{i+1} = e^{-\lambda_{T,a}\Delta t_i}\left(T^{(a)}_i + \lambda_{T,a}\Delta X_i\right), \qquad a \in \{1,2\}.$$

At each date, the conditional mean of the next return is estimated by ridge regression:

$$\widehat{B}_i = \arg\min_B \sum_{p=1}^{P} \left\| \frac{\Delta X^p_i}{\Delta t_i} - B^\top q^{\mathrm{ref}}_i(X^p_{1:i}) \right\|^2 + \rho_b \|B\|^2_F.$$

Hence,

$$\boxed{m^{\mathrm{ref}}_i(\boldsymbol{x}_{1:i}) = \widehat{B}_i^\top q^{\mathrm{ref}}_i(\boldsymbol{x}_{1:i})} \qquad \text{and} \qquad e^p_i = \Delta X^p_i - m^{\mathrm{ref}}_i(X^p_{1:i})\Delta t_i.$$

The reference drift used in the forward dynamics is $b^{\mathrm{ref}}_i = m^{\mathrm{ref}}_i$.

Construction of reference maps b^{ref} and σ^{ref} (2/2)

Reference marginal volatilities.
We maintain short- and long-memory trend and activity signals $T_i^{(a)}, V_i^{(a)} \in \mathbb{R}^d$, $a \in \{1,2\}$:

$$T_i = (1-\theta_T)T_i^{(1)} + \theta_T T_i^{(2)}, \qquad V_i = (1-\theta_V)V_i^{(1)} + \theta_V V_i^{(2)}.$$

The marginal volatility vector is

$$\boxed{v_i^{\mathrm{ref}} = \beta_0 + \beta_1 \odot T_i + \beta_2 \odot \sqrt{V_i}.}$$

The activity memories evolve as

$$V_{i+1}^{(a)} = e^{-\lambda_{V,a}\Delta t_i}\left(V_i^{(a)} + \lambda_{V,a}(v_i^{\mathrm{ref}})^{\odot 2}\Delta t_i\right).$$

To model cross-sectional dependence,

$$D_i^{\mathrm{ref}} = \mathrm{Diag}(v_i^{\mathrm{ref}}), \qquad C_i^{\mathrm{ref}} = D_i^{\mathrm{ref}} R D_i^{\mathrm{ref}},$$

where $R \in \mathbb{S}_{++}^d$ is a correlation matrix. Finally,

Calibration.
The volatility parameters and R are estimated from the residuals

$$e_i^p = \Delta X_i^p - m_i^{\mathrm{ref}}(X_{1:i}^p)\Delta t_i$$

by minimizing the Gaussian negative log-likelihood

$$\mathcal{L}_{\mathrm{ref}} = \frac{1}{P(N-1)} \sum_{p,i} \left[\log\det \sigma_i^{\mathrm{ref},p} + \frac{(e_i^p)^\top (C_i^{\mathrm{ref},p})^{-1} e_i^p}{2\Delta t_i} \right].$$

- Parameters are fitted on 80% of the training paths and selected on the remaining 20%.

Optimality system

The optimal control is characterized by the following forward-backward system, for $k \in [\![0, M-1]\!]$:

$$
\begin{cases}
X^\star_{s_{k+1}} = X^\star_{s_k} + b^{\mathrm{ref}}_{\eta(s_k)}\left(\boldsymbol{X}^\star_{1:\eta(s_k)}\right)\Delta s_k + \sigma^\star_{s_k}\sqrt{\Delta s_k}\,\epsilon_{s_{k+1}}, \\[2ex]
A^\star_{s_k} = \displaystyle\sum_{\substack{1\le j\le N\\ t_j\ge s_k}} \Bigg[\sum_{\substack{0\le l\le M-1\\ s_l\ge t_j}} \Big(\partial_{x_j} f_{\eta(s_l)}\big(\boldsymbol{X}^\star_{1:\eta(s_l)}, \sigma^\star_{s_l}\big) \\
\qquad\qquad + \partial_{x_j} b^{\mathrm{ref}}_{\eta(s_l)}\big(\boldsymbol{X}^\star_{1:\eta(s_l)}\big) A^\star_{s_{l+1}} \Big)\Delta s_l + \Gamma^G_j\big(\boldsymbol{X}^\star_{1:N}, \hat{\mu}^P\big) \Bigg], \\[2ex]
A^\star_{s_M} = \Gamma^G_N\big(\boldsymbol{X}^\star_{1:N}, \hat{\mu}^P\big), \\[2ex]
Z^\star_{s_k} = \mathbb{E}_{s_k}\left[A^\star_{s_{k+1}}\epsilon^\top_{s_{k+1}}\right], \\[2ex]
\sigma^\star_{s_k} = \left(\left(\sigma^{\mathrm{ref}}_{\eta(s_k)}\right)^{-1} + \dfrac{1}{\sqrt{\Delta s_k}} Z^\star_{s_k}\right)^{-1}.
\end{cases}
$$

For $(\boldsymbol{x}, \mu) \in (\mathbb{R}^d)^N \times \mathcal{P}_2((\mathbb{R}^d)^N)$, $\hat{\mu}^P := \frac{1}{P}\sum_{p=1}^P \delta_{\boldsymbol{X}^{(p)}_{1:N}}$ and $j \in [\![1, N]\!]$,

$$
\Gamma^G_j(\boldsymbol{x}, \mu) := \partial_{x_j}\frac{\delta}{\delta m} G(\mu \mid \nu)(\boldsymbol{x}) = w_{\mathrm{path}}\int_{(\mathbb{R}^d)^N} \partial_{x_j} k(\boldsymbol{x}, \boldsymbol{y})(\mu-\nu)(\mathrm{d}\boldsymbol{y})
$$
$$
+ w_{\mathrm{vol}}\int_{(\mathbb{R}^d)^N} \partial_{x_j} k(\varphi_{\mathrm{vol}}(\boldsymbol{x}), \varphi_{\mathrm{vol}}(\boldsymbol{y}))\,(\mu-\nu)(\mathrm{d}\boldsymbol{y}).
$$

We approximate $Z^\star$ by a recurrent neural network GRU_θ and iteratively update the forward process.

1. **Forward simulation.** For each iteration m and path p, initialize $X_{s_0}^{(m),p} = \xi$ and $h_{s_{-1}}^{(m),p} = 0$.

$$\left((\hat{Z}^{\theta^{(m)}})_{s_k}^{p}, h_{s_k}^{(m),p} \right) = \mathrm{GRU}_{\theta^{(m)}} \left(X_{s_k}^{(m),p}, h_{s_{k-1}}^{(m),p} \right),$$

$$\sigma_{s_k}^{(m),p} = \left((\sigma_{\eta(s_k)}^{\mathrm{ref}})^{-1} + \frac{1}{\sqrt{\Delta s_k}} (\hat{Z}^{\theta^{(m)}})_{s_k}^{p} \right)^{-1},$$

$$X_{s_{k+1}}^{(m),p} = X_{s_k}^{(m),p} + b_{\eta(s_k)}^{\mathrm{ref}} \left(\boldsymbol{X}_{1:\eta(s_k)}^{(m),p} \right) \Delta s_k + \sigma_{s_k}^{(m),p} \sqrt{\Delta s_k}\, \epsilon_{s_{k+1}}^{(m),p}.$$

The hidden state $h_{s_k} \in \mathbb{R}^{96}$ provides a learned summary of the path history.

$$\mathrm{GRU}_{\theta^{(m)}} \longrightarrow Z^{(m)} \longrightarrow \sigma^{(m)} \longrightarrow X^{(m)}$$

2. Ridge proxy.
The pathwise adjoint $A^{(m),p}$ is computed following previous relation.

$$u_{s_k}^{(m),p} = \Phi_k^{\mathrm{ref}}\left(\boldsymbol{X}_{\mathbf{1}:s_k}^{(m),p}\right).$$

We regress the realized adjoint observations $A_{s_{k+\mathbf{1}}}^{(m),p}\epsilon_{s_{k+\mathbf{1}}}^{(m),p}$:

$$\beta_{s_k}^{(m+\mathbf{1})} = \arg\min_{\beta}\left\{\sum_{p=\mathbf{1}}^{P}\left|A_{s_{k+\mathbf{1}}}^{(m),p}\epsilon_{s_{k+\mathbf{1}}}^{(m),p} - \beta^\top u_{s_k}^{(m),p}\right|^{\mathbf{2}} + \gamma|\beta|^{\mathbf{2}}\right\}.$$

Hence,

$$Z_{s_k}^{\mathrm{proxy},(m+\mathbf{1}),p} = (\beta_{s_k}^{(m+\mathbf{1})})^\top u_{s_k}^{(m),p}.$$

3. Neural projection.
The GRU is trained to reproduce the ridge proxy:

$$\mathcal{L}_{\mathrm{adj}}^{(m+\mathbf{1})}(\theta) = \frac{\mathbf{1}}{PM}\sum_{p=\mathbf{1}}^{P}\sum_{k=\mathbf{0}}^{M-\mathbf{1}}\left\|(\hat{Z}^\theta)_{s_k}^p - Z_{s_k}^{\mathrm{proxy},(m+\mathbf{1}),p}\right\|^{\mathbf{2}}.$$

Then,

$$\theta^{(m+\mathbf{1})} = \arg\min_{\theta\in\Theta}\mathcal{L}_{\mathrm{adj}}^{(m+\mathbf{1})}(\theta).$$

The procedure is iterated for $m = 0, \ldots, K-1$.

$$X^{(m)} \to A^{(m)} \to Z^{\mathrm{proxy},(m+1)} \to \theta^{(m+1)} \to X^{(m+1)}$$